

A Dempster-Shafer Theoretic Approach to Understanding Indirect Speech Acts

Tom Williams¹, Rafael C. Núñez², Gordon Briggs¹,
Matthias Scheutz¹, Kamal Premaratne², and Manohar N. Murthi²

¹ Human-Robot Interaction Laboratory, Tufts University
200 Boston Avenue, Medford MA 02155, USA
{williams,gbriggs,mscheutz}@cs.tufts.edu

² Department of Electrical and Computer Engineering, University of Miami
1251 Memorial Drive, Coral Gables, FL 33146, USA
nunez@miami.edu, {kamal,mmurthi}@miami.edu

Abstract. Understanding *Indirect Speech Acts (ISAs)* is an integral function of human understanding of natural language. Recent attempts at understanding ISAs have used rule-based approaches to map utterances to deep semantics. While these approaches have been successful in handling a wide range of ISAs, they do not take into account the uncertainty associated with the utterance’s context, or the utterance itself. We present a new approach for understanding ISAs using the Dempster-Shafer theory of evidence and show how this approach increases the robustness of ISA inference by (1) accounting for uncertain implication rules and context, (2) fluidly adapting rules given new information, and (3) enabling better modeling of the beliefs of other agents.

Key words: Speech act theory, intention understanding, Dempster-Shafer theory

1 Introduction

People do not often directly express their intentions. For various social reasons (e.g. politeness [2]) they instead use linguistic strategies such as *Indirect Speech Acts (ISAs)*, i.e., utterances whose intended meanings differ from their literal meanings. Two approaches have been proposed for handling ISAs in computational systems. The *inferential* approach reasons about possible intended meanings by considering observed speech acts in the current context as part of a broader plan that captures the agent’s goals and intentions [9]. In contrast, the *idiomatic* approach leverages the fact that certain ISA forms are *conventionalized*, i.e., that they are directly associated with an inferred meaning, largely independent of context [15]. Both approaches have advantages and disadvantages: the inferential approach requires the ability to infer interlocutors’ possible plans, which can be quite computationally expensive, while the idiomatic approach is limited to ISAs for which conventionalized meanings exist. We contend that there are three capabilities necessary for robust understanding of conventionalized ISAs through the idiomatic approach:

C1: Uncertainty. An agent must not assume perfect knowledge of the contexts in which an indirect interpretation applies. The conventionalized meaning of an ISA is not always the intended meaning; sometimes “I’d love some cake” is simply a statement expressing a desire, and not an indirect request for someone to give you cake. Since an agent might not always be able to determine the true intended meaning of an utterance, it should ascribe a level of confidence to each of its interpretations, based on the contextual factors that provide evidence for each interpretation. Furthermore, it is important that an agent be able to represent and reason about its own uncertainty and ignorance, and be able to act appropriately when uncertainty is identified.

C2: Adaptation. Since an agent should be able to learn new ISAs, and since it may not know the precise scenarios in which new ISAs should be used, an agent should be able to learn and adapt new rules, using feedback from interlocutors to adjust its beliefs as to when the rules it knows apply. For example, consider the following dialogue:

DATA: Are you certain you do not wish to talk about your mother?

GEORDI: Why do you ask that?

DATA: You are no doubt feeling emotional distress as a result of her disappearance. Though you claimed to be “just passing by,” that is most likely an excuse to start a conversation about this uncomfortable subject. Am I correct?

GEORDI: Well, no. Sometimes “just passing by” means “just passing by.”

DATA: Then I apologize for my premature assumption. . .

GEORDI: Data, maybe you gave up too fast.

DATA: I do not understand.

GEORDI: When I said “just passing by” means “just passing by,” I didn’t really mean it.

DATA: My initial assumption was correct. You do wish to speak of your mother.

Short dialogue from Star Trek: The Next Generation. “Interface”

In the space of this short dialogue, an agent (i.e., the android “Data”) must make several adaptations. First, he must alter his beliefs about the ISA “just passing by” based on feedback from Geordi that the ISA’s literal meaning had been the correct interpretation. Then, he must at least partially revert to his previous beliefs, as well as alter his belief as to when $said(X, Y) \rightarrow means(X, Y)$.

C3: Belief modeling. An agent should be able to model interlocutors’ beliefs: the interpretation of an ISA *uttered by an interlocutor* should be based not on the *robot’s* beliefs about, for example, its capabilities and obligations, but rather on its *interlocutor’s* beliefs.

In this paper, we present a novel approach that addresses the three aforementioned capabilities to robustly handle idiomatic ISAs. We first give a brief overview of Dempster-Shafer (DS) theory and DS rule-based inference before presenting the DS theoretic approach to ISA understanding. We then compare this approach to previous work, and conclude with an outlook for future work.

2 Indirect Speech Act (ISA) Modeling

Suppose a robot were told “I would love a coffee.” This utterance was likely generated due to some intention to communicate something to the robot. This intention was likely formed due to some contextual factors, whether environmental (e.g., the interlocutor was tired) or dialogic (e.g., the robot had just asked the interlocutor “Would you like a coffee?”).

We see two distinct ways to model these contextual effects. **(A)** They could be modeled as part of the user’s intentions; the relationship between the robot and its interlocutor may, e.g., cause the interlocutor to form the intention to avoid production of utterances that could be insulting to the robot. **(B)** Alternatively, context could directly affect utterance production, and the information the interlocutor intends to communicate as well as the context that dictates whether and how the information is conveyed could be treated as separate factors leading to the production of the utterance. We have chosen approach **(B)**, using a graphical model with random variables C , U and I such that I depends on C and U depends on C and I . Here, C , I and U represent distributions over possible contexts, intentions and utterances, respectively. From this model, we are interested in inferring the interlocutor’s intentions given the current context and the produced utterance. The Bayesian approach to this inference problem would be to calculate $P(U, I, C) = P(U|I, C) \cdot P(I|C) \cdot P(C)$. To calculate $P(I|U, C)$ given an utterance u and context c , one would then formulate:

$$P(I|U = u, C = c) = \frac{P(U = u|I, C = c) \cdot P(I|C = c) \cdot P(C = c)}{\sum_{i \in I} P(U = u, I = i, C = c)}.$$

However, $P(U|I, C)$ is at least as hard to calculate as $P(I|U, C)$, for two reasons. First, we do not have access to the distribution over an interlocutor’s intentions as we cannot look inside his or her head. Second, one would need a table containing priors on all combinations of intentions and contexts; a table that could not be realistically represented unless sparse representations were used. Even if such a table could be constructed, it is unclear where its values would come from. An example of a Bayesian approach to utterance interpretation can be found in [3]. However, this work appears to only engage in speech act classification and not semantic analysis of utterances.

3 DS-Based Inference for ISAs

Because the direct Bayesian approach of inferring $P(I|U, C)$ by way of $P(U|I, C)$, $P(I|C)$, and $P(C)$ (i.e., the conditional probability of utterances occurring given intentions and context, the conditional probability of intentions given context, and the prior distribution of contexts) does not make the inference problem any easier, we instead tackle $P(I|U, C)$ directly. To do so, we create rules of the form $u \wedge c \Rightarrow_{[\alpha, \beta]} i$. Here, u is an utterance, c is a context, i is an intention, and $[\alpha, \beta]$, where $0 \leq \alpha \leq \beta \leq 1$, is the uncertainty interval associated with the rule.

3.1 Dempster-Shafer (DS) Theory

DS theory is an uncertainty processing framework often interpreted as an extension of the Bayesian framework [13,8]. Its notions of belief and plausibility bear a close relationship to the inner and outer measures in probability theory [4].

Basic Notions in DS Theory

Frame of Discernment. In DS theory, the discrete set of elementary events of interest related to a given problem is called the *Frame of Discernment (FoD)*. We take the FoD to be the finite set of mutually exclusive events $\Theta = \{\theta_1, \theta_2, \dots, \theta_N\}$. The power set of Θ is denoted by $2^\Theta = \{A : A \subseteq \Theta\}$.

Basic Belief Assignment. A *Basic Belief Assignment (BBA)* is a mapping $m_\Theta(\cdot) : 2^\Theta \rightarrow [0, 1]$ such that $\sum_{A \subseteq \Theta} m_\Theta(A) = 1$ and $m_\Theta(\emptyset) = 0$. The BBA measures the support assigned to propositions $A \subseteq \Theta$ only. The subsets of A with nonzero mass are referred to as *focal elements*, and comprise the *core* \mathcal{F}_Θ . The triple $\mathcal{E} = \{\Theta, \mathcal{F}_\Theta, m_\Theta(\cdot)\}$ is called the *Body of Evidence (BoE)*.

Belief, Plausibility, and Uncertainty. Given a BoE $\{\Theta, \mathcal{F}, m\}$, the *belief* $\text{Bl} : 2^\Theta \rightarrow [0, 1]$ is $\text{Bl}(A) = \sum_{B \subseteq A} m_\Theta(B)$. So, $\text{Bl}(A)$ captures the total support that can be committed to A without also committing it to the complement A^c of A . The *plausibility* $\text{Pl} : 2^\Theta \rightarrow [0, 1]$ is $\text{Pl}(A) = 1 - \text{Bl}(A^c)$. So, $\text{Pl}(A)$ corresponds to the total belief that does not contradict A . The *uncertainty* of A is $[\text{Bl}(A), \text{Pl}(A)]$.

Conditional Fusion Equation (CFE). The evidence from two sources having the BBAs $m_j(\cdot)$ and $m_k(\cdot)$ can be fused using various fusion strategies. A robust fusion strategy is the *Conditional Fusion Equation (CFE)* [11].

Uncertain Logic. Uncertain logic, a DS-based extension of classical logic which deals with propositions whose truth is uncertain, handles expressions of the following form [7,6]:

$$\varphi(x), \text{ with uncertainty } [\alpha, \beta], 0 \leq \alpha \leq \beta \leq 1, \quad (1)$$

where $\varphi(x)$ is a proposition which contains a reference to individual x , $x \in \mathcal{D} = \{x_1, x_2, \dots, x_n\}$, a finite set of individuals. A DS model for expression (1) can be defined over the logical FoD $\Theta_{\varphi, x} = \{\varphi(x) \times \mathbf{1}, \varphi(x) \times \mathbf{0}\}$, which contains two mutually exclusive elements: our confidence that the proposition φ applies and does not apply to x , respectively. When no confusion can arise, we represent these two elements as $\{x, \bar{x}\}$. Then the information in (1) can be captured by the following DS model over $\Theta_{\varphi, x}$: $m(x) = \alpha$; $m(\bar{x}) = 1 - \beta$; $m(\Theta_{\varphi, x}) = \beta - \alpha$. In general, we could also model the uncertainty of propositions $\varphi_i \in \{\varphi_1, \dots, \varphi_M\}$ applying to particular elements $x_j \in \Theta_x$, i.e.,

$$\varphi_i(x_j), \text{ with uncertainty } [\alpha_{i,j}, \beta_{i,j}], 0 \leq \alpha_{i,j} \leq \beta_{i,j} \leq 1, \quad (2)$$

via a model defined over the FoD $\Theta_{\varphi_i, x_j} = \{\varphi_i(x_j) \times \mathbf{1}, \varphi_i(x_j) \times \mathbf{0}\} = \{x_{i,j}, \bar{x}_{i,j}\}$.

CFE-Based Logical Operators: We can now define logic operations such as NOT (\neg), AND (\wedge), and OR (\vee) [7,6]. Whenever possible we define operations in a simple unquantified first-order logic model (e.g., based on (1)) instead of (2), but extension to more complex cases and simplification into propositional logic are straightforward.

Logical Negation. Logical negation of uncertain proposition $\varphi(x)$ in (1) and its corresponding DS model are $\neg\varphi(x)$, with uncertainty $[1 - \beta, 1 - \alpha]$, and

$$m(x) = 1 - \beta; \quad m(\bar{x}) = \alpha; \quad m(\Theta_{\varphi, x}) = \beta - \alpha. \quad (3)$$

Logical AND/OR. Consider M logic predicates, each providing a statement regarding the truth of x with respect to the proposition $\varphi_i(\cdot)$ in (2). Then, the corresponding DS models for $\varphi_i(x)$ are, for $i = 1, 2, \dots, M$:

$$m_i(x) = \alpha_i; \quad m_i(\bar{x}) = 1 - \beta_i; \quad m_i(\Theta_{\varphi_i, x}) = \beta_i - \alpha_i. \quad (4)$$

The DS model for the logical AND and OR of the statements in (4) can be defined as:

$$m_{\wedge}(\cdot) = \bigcap_{i=1}^M m_i(\cdot); \quad m_{\vee}(\cdot) = \left(\bigcap_{i=1}^M m_i^c(\cdot) \right)^c, \quad (5)$$

respectively, where \bigcap denotes an appropriate fusion operator. When the CFE (with appropriate parameters) is used as the fusion operator, consistency with classical logic can be achieved [7]. In the case of CFE-Based Logical Operators, logical AND when $M = 2$ is defined as:

$$m(x) = \underline{\alpha}; \quad m(\bar{x}) = 1 - \underline{\beta}; \quad m(\Theta_{\varphi_1, x} \times \Theta_{\varphi_2, x}) = \underline{\beta} - \underline{\alpha}, \quad (6)$$

where $\underline{\alpha} = \min(\alpha_1, \alpha_2)$ and $\underline{\beta} = \min(\beta_1, \beta_2)$, and logical OR is defined as:

$$m(x) = \bar{\alpha}; \quad m(\bar{x}) = 1 - \bar{\beta}; \quad m(\Theta_{\varphi_1, x} \times \Theta_{\varphi_2, x}) = \bar{\beta} - \bar{\alpha}, \quad (7)$$

where $\bar{\alpha} = \max(\alpha_1, \alpha_2)$ and $\bar{\beta} = \max(\beta_1, \beta_2)$. Hereafter, $m_1 \otimes m_2$ denotes the DS model corresponding to the uncertain logic operation $\varphi_1(\cdot) \wedge \varphi_2(\cdot)$.

Logical Implication. Given two logic statements $\varphi_1(\cdot)$ and $\varphi_2(\cdot)$, an implication rule in classical logic takes the form $\varphi_1(x) \Rightarrow \varphi_2(y) \equiv \neg\varphi_1(x) \vee \varphi_2(y) \equiv \neg(\varphi_1(x) \wedge \neg\varphi_2(y))$, where $x_i \in \Theta_x$ and $y_j \in \Theta_y$.

As shown in [6], the DS model for the uncertain implication $\varphi_1(\cdot) \Rightarrow \varphi_2(\cdot)$ over the true-false FoD $\{\mathbf{1}, \mathbf{0}\}$ may be defined using (7) and (3):

$$\begin{aligned} m_{\varphi_1 \rightarrow \varphi_2}(\mathbf{1}) &= \alpha_R; \quad m_{\varphi_1 \rightarrow \varphi_2}(\mathbf{0}) = 1 - \beta_R; \\ m_{\varphi_1 \rightarrow \varphi_2}(\{\mathbf{1}, \mathbf{0}\}) &= \beta_R - \alpha_R, \end{aligned} \quad (8)$$

where $\alpha_R = \max(1 - \beta_1, \alpha_2)$ and $\beta_R = \max(1 - \alpha_1, \beta_2)$. Thus, the implication rule's uncertainty interval is $[\alpha_R, \beta_R]$. This DS model provides us with an important inference tool. Suppose DS models for the implication rule and antecedent are known. We then obtain the following DS model for the consequent [6]:

$$\alpha_2 = \begin{cases} \alpha_R, & \text{if } \alpha_R > 1 - \beta_1; \\ 0, & \text{if } \alpha_R = 1 - \beta_1; \\ \text{no solution,} & \text{otherwise;} \end{cases} \quad \text{and } \beta_2 = \begin{cases} \beta_R, & \text{if } \beta_R > 1 - \alpha_1; \\ \text{no solution,} & \text{otherwise.} \end{cases} \quad (9)$$

Inference. Inference in uncertain logic shares the fundamental principles of classical logic, and adds the possibility of attaching, tracking, and propagating uncertainties that may arise on premises and/or rules. The model in (9) can be used as an uncertain Modus Ponens (MP) rule [7]. We use $m_2 = m_1 \odot m_{12}$ to express that the BBA m_2 is obtained after applying MP when the BBAs of the antecedent m_1 and the implication $m_{12} = m_{\varphi_1 \rightarrow \varphi_2}$ are known.

Symbolic Dempster-Shafer Operators: In [14], Tang et al. produce another candidate set of operators for Logical AND, OR, Implication, and Modus Ponens (their operator for logical negation is equivalent to that defined by Núñez et al). **Logical AND/OR.** Tang et al. define the DS model for logical AND as:

$$\begin{aligned} m(x) &= \alpha_1 * \alpha_2; \quad m(\bar{x}) = 1 - (\beta_1 * \beta_2); \\ m(\Theta_{\varphi_1, x} \times \Theta_{\varphi_2, x}) &= (1 - m(\bar{x})) - m(x), \end{aligned} \quad (10)$$

and OR as:

$$\begin{aligned} m(x) &= \frac{\alpha_1 + \alpha_2}{2}; \quad m(\bar{x}) = \frac{(1 - \beta_1) + (1 - \beta_2)}{2}; \\ m(\Theta_{\varphi_1, x} \times \Theta_{\varphi_2, x}) &= (1 - m(\bar{x})) - m(x), \end{aligned} \quad (11)$$

Logical Implication. Tang et al. define logical implication as

$$\begin{aligned} m(x) &= \frac{(1 - \alpha_1) + \alpha_2}{2}; & m(\bar{x}) &= \frac{\beta_1 + (1 - \beta_2)}{2}; \\ m(\Theta_{\varphi_1, x} \times \Theta_{\varphi_2, x}) &= (1 - m(\bar{x})) - m(x), \end{aligned} \quad (12)$$

allowing the consequent to be calculated as

$$\alpha_2 = \alpha_1 * \alpha_R; \quad \beta_2 = \beta_R. \quad (13)$$

This as well can also be used as an uncertain Modus Ponens rule [14].

3.2 Inferring Intentions

Let $\Theta_U = \{u_1, u_2, \dots, u_{N_u}\}$ be the set of all utterances an agent may interpret, and $\Theta_C = \{c_1, c_2, \dots, c_{N_c}\}$ be the set of all contextual items. Also, let $\Theta_I = \{i_1, i_2, \dots, i_{N_i}\}$ be the set of atomic intentions an interlocutor may be trying to communicate. Using the uncertain logic framework described above, we can define the BBAs $m_u(\cdot), m_c(\cdot), m_i(\cdot)$ over the FoDs $\Theta_U, \Theta_C, \Theta_I \times \{\mathbf{1}, \mathbf{0}\}$ respectively. The information required to calculate $m_i(\cdot)$ is available to the agent: its natural language understanding system can provide a distribution over possible utterances heard (yielding m_u), its knowledge base can provide a distribution over different contextual items being believed (yielding m_c), and information regarding the uncertainty of $(u \wedge c) \Rightarrow i$ can be encoded in a table M indexed by utterance u and contextual item c , defining a BBA $m_{uc \rightarrow i}$. Using these three BBAs, we can obtain a model for the uncertainty of intention i through MP by computing $m_i(\cdot) = ((m_u \otimes m_c) \odot m_{uc \rightarrow i})(\cdot)$ defined over the FoD $\Theta_I \times \{\mathbf{1}, \mathbf{0}\}$.

While recent approaches to ISA understanding (e.g., [1]) use only the first applicable rule found, we instead combine multiple applicable rules, as multiple rules may produce different beliefs regarding the same hypothesis. For example, one rule may produce an inference that an interlocutor does not want coffee as he rarely drinks it, while another may produce an inference that he does because he just stated how tired he was. Fusing these results yields a single inference that paints a better picture of the agent's confidence. We can obtain a DS model produced by n applicable rules of the type $(u \wedge c) \Rightarrow i$ by combining n BBAs:

$$m_\psi(\cdot) = \bigcap_{u \in \Theta_U, c \in \Theta_C} ((m_u \otimes m_c) \odot m_{uc \rightarrow i})(\cdot), \quad (14)$$

where \cap refers to some generic fusion operator (e.g., the CFE). Since each BBA resulting from the application of MP in the equation above is defined over the FoD $\Theta_I \times \{\mathbf{1}, \mathbf{0}\}$, so is the resulting fused DS model m_ψ .

3.3 Algorithm

Given the BoEs $\{\Theta_U, m_u\}$ and $\{\Theta_C, m_c\}$, which encode the uncertainty as to the truthfulness of utterance u and context c respectively, and a list of applicable rules R , Algorithm 1 infers the intended meaning of u . The algorithm first collects the consequents resulting from the application of u and c to rule r into the set S , and then groups these consequents using $group(S)$ such that the consequents in each group are all on the same

Algorithm 1 $\text{getIntendedMeaning}(\{\Theta_U, m_u\}, \{\Theta_C, m_c\}, R)$

```

1:  $\{\Theta_U, m_u\}$ : BoE of candidate utterances
2:  $\{\Theta_C, m_c\}$ : BoE of relevant contextual items
3:  $R$ : Currently applicable rules
4:  $S = \emptyset$ 
5: for all  $r \in R$  do
6:    $S = S \cup \{(m_u \otimes m_c) \odot m_{r=uc \rightarrow i}\}$ 
7: end for
8:  $G = \text{group}(S)$ 
9:  $\psi = \emptyset$ 
10: for all group  $g_i \in G$  do
11:    $\psi = \psi \cup \{\bigcap_{j=0}^{|g_i|} g_{i,j}\}$ 
12: end for
13: return  $\psi$ 

```

FoD (for example, a consequent $wants(jim, coffee)[0.1, 0.3]$ would be in the same group as $wants(jim, coffee)[0.7, 0.9]$ but in a different group than $wants(jim, tea)[0.4, 0.6]$. This allows the constituents of each group to be fused using the CFE, or to be processed using MP. Finally, the fused consequents resulting from this step are collected into ψ , which is returned.

Our approach is similar to that of [1], who express each rule as a tuple $(\tilde{C}, \tilde{U}, [[U]]_C)$ where \tilde{C} is a set of contextual constraints, \tilde{U} is an utterance form, and $[[U]]_C$ is a set of belief updates to be made if \tilde{U} and \tilde{C} match the current utterance and context. Rules are sequentially compared against utterances and contexts. If a matching rule is found, its consequent is immediately returned. As only one rule is ever applied, specific rules are written for particular combinations of contextual items, and are arranged in descending order of specificity. This differs from the DS-theoretic approach, in which the consequents from all applicable rules are combined; instead of a single rule encoding all pieces of context that evidence a given intention, multiple rules are used.

4 Evaluation

We will now present an evaluation of our algorithm and demonstrate the capabilities facilitated by our approach. The evaluation of a system at this stage of the natural language pipeline is difficult, as the performance of the algorithm is tightly coupled with the performance of components that precede it in the natural language pipeline (e.g., speech recognition, parsing, semantic analysis). We thus take the same approach to evaluation as previous work, i.e., through a case study that demonstrates the behavior of our algorithm. We will now show how our algorithm works towards the capabilities necessary for robust understanding of conventionalized ISAs, and then compare our algorithm to previous work.

4.1 Handling Uncertainty

Consider a robot speaking with interlocutor Jim. Suppose Jim says to the robot: "I need coffee". From the robot's perspective, this utterance is represented as

VIII

$Stmt(jim, self, needs(jim, coffee))$). In this representation, the three arguments represent the speaker of the utterance (in this case, Jim), the receiver (in this case, the robot ("self")), and the conveyed message (in this case, *Jim needs coffee*). Suppose the robot knows two pragmatic rules. In both rules, $[\alpha_{R_i}, \beta_{R_i}]$ represents the belief and plausibility of rule i .

First, if B believes A is a barista, then telling A that B needs coffee indicates that B wants A to believe that B wants A to get them coffee.

$$r_{[\alpha_{R_0}, \beta_{R_0}]}^0 = \frac{\begin{array}{l} (Context : believe(B, barista(A))) \\ (Utterance : Stmt(B, A, need(A, coffee))) \end{array}}{(Intention : want(B, believe(A, want(B, get_for(A, B, coffee))))}$$

Second, if B believes C is thirst quenching, telling A that B needs C indicates that B wants A to believe that B is thirsty.

$$r_{[\alpha_{R_1}, \beta_{R_1}]}^1 = \frac{\begin{array}{l} (Context : believe(B, quenches(C, thirst))) \\ (Utterance : Stmt(B, A, need(A, C))) \end{array}}{(Intention : want(B, believe(A, thirsty(B))))}$$

Our approach affords the first capability of an ideal system, i.e., the ability to handle uncertain contextual and dialogical information, and to recognize and reason about one's own ignorance. To demonstrate this, suppose the robot strongly believes the following:

(a) Jim believes coffee is thirst quenching: $bel(jim, quenches(coffee, thirst))$ [1.0, 1.0], **(b)** Jim just said he needs a coffee: $Stmt(jim, self, need(jim, coffee))$ [0.9, 0.9], and **(c)** Jim may or may not think the robot is a barista: $bel(jim, barista(self))$ $[\alpha_b, \beta_b]$.

Applying rules r^0 and r^1 will produce a BBA that encompasses information from two consequents, namely:

$$\begin{aligned} c_0 &= want(jim, bel(self, want(jim, get_for(self, jim, coffee)))) \\ c_1 &= want(jim, bel(self, thirsty(jim))). \end{aligned}$$

The degree to which c_0 and c_1 are believed depends on the logic models and operators that are being used. In our example, we use Núñez and Tang's operators (see Section 3.1 above). Although Núñez' logic models can be parameterized to enforce certain logic properties, for ease of explanation we only consider the CFE-based classically consistent uncertain logic operators defined in [7]. Table 1 contains some cases that illustrate how the uncertainty in the two consequents changes depending on the set of logic models and fusion operators, the degree to which the robot believes the interlocutor believes the robot is a Barista (b), and the degree to which the robot believes the two rules r^0 and r^1 hold.

Note that our approach can modulate its interpretation of utterances based on the certainty of the relevant utterance, contextual factors, and pragmatic rules. However, an ideal system should also explicitly reason about its own ignorance.

Table 1. Comparison of operators under Tang and Núñez

	$b[\alpha, \beta]$	$r^0[\alpha, \beta]$	$r^1[\alpha, \beta]$	Fusion	$c_0[\alpha, \beta]$	$c_1[\alpha, \beta]$	λ_0	λ_1
1	[0.9,0.9]	[0.85,0.9]	[0.7,0.85]	Núñez	[0.85,0.90]	[0.70,0.85]	0.41	0.17
2	[0.9,0.9]	[0.85,0.9]	[0.7,0.85]	Tang	[0.69,0.90]	[0.63,0.85]	0.18	0.11
3	[0.1,0.1]	[0.1,0.1]	[0.5,0.5]	Núñez	N/A	[0.50,0.50]	N/A	0.00
4	[0.1,0.1]	[0.1,0.1]	[0.5,0.5]	Tang	[0.01,0.10]	[0.05,0.50]	0.56	0.07
5	[0.5,0.5]	[0.1,0.5]	[0.5,0.5]	Núñez	N/A	N/A	N/A	N/A
6	[0.5,0.5]	[0.1,0.5]	[0.5,0.5]	Tang	[0.05,0.5]	[0.45,0.5]	0.07	0.002
7	[0.002,0.002]	[0.99,0.99]	[0.99,0.99]	Núñez	[0.99,0.99]	[0.99,0.99]	0.92	0.92
8	[0.002,0.002]	[0.99,0.99]	[0.99,0.99]	Tang	[0.002,0.99]	[0.80,0.99]	0.00001	0.35

Since we are using a DS-theoretic approach, we can use the consequents' uncertainty intervals to determine whether or not the agent needs to ask for clarification. Specifically, we use the ambiguity measure defined in [6]:

$$\lambda = 1 + \frac{\beta}{1 + \beta - \alpha} \log_2 \frac{\beta}{1 + \beta - \alpha} + \frac{1 - \alpha}{1 + \beta - \alpha} \log_2 \frac{1 - \alpha}{1 + \beta - \alpha}.$$

For example, for the interval [0.6, 0.9], $\lambda = 1 + \frac{0.9}{1.3} \log_2 \frac{0.9}{1.3} + \frac{0.4}{1.3} \log_2 \frac{0.4}{1.3} = 0.11$. $\lambda \rightarrow 0$ as uncertainty grows and as α and $1 - \beta$ grow closer together. Using this equation, we generate a clarification request if $\lambda \leq 0.1$. This makes use of information that is unavailable to the Bayesian approach.

We will now briefly compare Tang and Núñez' fusion operators before discussing the other capabilities afforded by our approach. Referring to Table 1, one of the most visible results, is that there are several cases in which Núñez' operators do not return a solution (see rows 3 and 6). This is expected since, based on the CFE-based classically consistent logic models, a Modus Ponens does not return a result if there is not enough evidence in the antecedent that supports making a conclusion. In most cases, this lack of supporting evidence is shown as very low values for λ_0 and λ_1 when using Tang's operators, with the only exception being the case of λ_0 in row 4, where the very low uncertainty in the antecedent translates into a very low uncertainty in the consequent.

A second difference between Tang and Núñez' operators is that Tang's operators seem to be slightly more conservative in the allocation of evidence. This can be seen in rows 1 and 2, where the uncertainty intervals associated with c_0 and c_1 are wider when Tang's operators are used.

A more important difference between these operators is evidenced in rows 7 and 8. In the scenarios depicted in these rows, the application of Modus Ponens based on Núñez' operators leads to a potentially problematic high confidence in the consequents c_0 and c_1 . Note that, in row 7, the very small uncertainty associated with the antecedent (which suggests that it is false) is not reflected in the resulting uncertainty for the consequent c_0 . Furthermore, λ_0 renders a high value, preventing the automatic request of additional evidence by the robot. Due to the more conservative allocation of evidence of Tang's operators, this problem is not visible in row 8. In light of this, we prefer the use of Tang's logical operators in the domain of ISA understanding and pragmatic inference. Using

Núñez’ operators in this domain could be enabled by a different parameterization of the uncertain logic (e.g., with CFE coefficients that relax some classical logic properties), or by incorporating additional components in the reasoning system (e.g., additional logic rules) aimed at solving the above mentioned issue.

4.2 Adaptation

The second capability of an ideal system is the ability to adapt old rules and learn new ones. We currently assume that the initial beliefs and plausibilities of our rules and contextual items are given, but we do allow rules to be adapted based on user feedback. Upon receiving a corrected rule from a user, it is compared against all current rules. Those whose antecedents and consequents are on the same frames as the antecedents and consequents of the new rule may be updated using the Conditional Update Equation (CUE) as defined in [11]. For example, if rule r^i is on interval $[0.8, 0.8]$, and a correction states that in the current context, $[\alpha_{R_i}, \beta_{R_i}]^i$ should be $[0.5, 0.9]$, the CUE will update the rule’s uncertainty to $r^i_{[0.53, 1.0]}$ (a substantial increase in uncertainty). Although the proposed approach only allows for adaptation of rules, it could easily be extended to allow for the addition of new rules, which would initially have very high levels of uncertainty and would become less uncertain with exposure to applications of the rule.

4.3 Belief Modeling

The third capability of an ideal system is the ability to reason about other agents’ beliefs. Rules such as r^0 and r^1 are formulated in terms of the *interlocutor’s* beliefs; to determine what interlocutor J is trying to communicate, J ’s utterances must be evaluated in the context of J ’s beliefs. For example, if J says he needs coffee, the likelihood that he is trying to order a coffee should be modulated not by the *robot’s* belief that it is a barista, but instead by J ’s beliefs; if J has no reason to think the robot is a barista, his statement should not be viewed as a coffee order even if the robot has barista training. Belief modeling also allows natural representation of interlocutors’ beliefs about the robot’s abilities and social roles. For example, the robot may need general rules (e.g., Equation 15) that suggest that a statement such as “I need a coffee” is only an indirect request if its interlocutor believes the robot to be able and obligated to get them coffee.

$$\frac{\begin{array}{l} (\text{Context} : \text{bel}(B, \text{obligated}(A, \text{give}(A, B, C)))) \\ (\text{Utterance} : \text{Stmt}(B, A, \text{would like}(B, C))) \end{array}}{(\text{Intention} : \text{want}(B, \text{bel}(A, \text{want}(\underline{\text{give}}(A, B, C)))))} \quad (15)$$

4.4 Comparison to previous work

While ISAs have been studied for nearly forty years in philosophy and linguistics [12,10], few computational approaches have been presented for modeling idiomatic ISAs in situated contexts (e.g., [15,1]). We believe that the DS theoretic approach represents a significant advance over existing approaches.

Wilske and Kruijff’s proposal maps indirect requests to action requests [15], and models ambiguity and adaptation: certain utterance types and dialogue

Table 2. Comparison of new and existing work

	Wilske	Hinkelman	Briggs	DeVault	Proposed
Conventional ISAs	•	•	•		•
Unconventional ISAs		•	•		
Handles uncertain context					•
Handles uncertain utterances				•	•
Handles uncertain rules					•
Robust rule combination		•			•
Models agent’s ignorance	•				•
Adaptation of existing rules	•			•	•
Learning of new rules				•	
Uses belief modeling			•		•

contexts will prompt a clarification request whose result determines whether the agent will change its belief. However, such changes are all-or-nothing, which can lead to shifts in belief of unwarranted magnitude. Wilske and Kruijff attempt to rectify this problem by always allowing a chance for the agent to ask for clarification, so unwarranted belief shifts can be reversed. However, this can lead to superfluous questions (when the agent is fairly certain) and incorrect interpretations (when the agent has a belief that is certain and incorrect).

DeVault and Stone presented *COREF*, a dialog system that uses observed dialog features to learn the appropriate interpretations of utterances. While *COREF* learns to identify the appropriate meaning of an utterance, this consists of resolving lexical, referential, and dialog-move ambiguities in a simple shape-identification game; *COREF* does not show evidence of handling ISAs.

Some systems that handle conventionalized ISAs also handle unconventionalized ISAs using plan reasoning [5,1]. These approaches first attempt to handle conventionalized ISAs via rule-based systems that map incoming utterances to deep semantics according to context, and then handle unconventionalized ISAs using plan-reasoning. These approaches do not handle uncertainty or adaptation.

5 Conclusion

We have presented a novel approach for robustly handling ISAs using DS-theoretic uncertain logical inference, and have shown (1) how the proposed algorithm robustly deals with uncertainty in implication rules and dialogic and environmental context, (2) how belief modeling allows the algorithm to better resolve ISAs, and (3) how rules can be adapted. Table 2 demonstrates that this approach comes closer than previous approaches to satisfying these capabilities.

Currently, our algorithm only handles conventionalized ISAs. A logical next step is to use a hierarchical approach like that described by [1]. While Briggs et al. attempt to understand an ISA inferentially only if a conventionalized form does not exist, we would also need to attempt to understand ISAs inferentially if idiomatic analysis only produced consequents with very low belief and/or high uncertainty. It would also be advantageous to extend the adaptation algorithm

to learn *new* ISAs when a correction is provided for which an existing rule does not exist. Future work includes determining which set of fusion operators is preferable, and then performing a large scale evaluation of performance under that set of fusion operators when integrated into a cognitive robotic architecture.

6 Acknowledgments

This work was funded in part by ONR grants #N00014-11-1-0493 and #N00014-10-1-0140, and in part by NSF grants #1111323 and #1038257.

References

1. Briggs, G., Scheutz, M.: A hybrid architectural approach to understanding and appropriately generating indirect speech acts. In: Proceedings of the 27th AAAI Conference on Artificial Intelligence (July 2013)
2. Brown, P.: *Politeness: Some Universals in Language Usage*, vol. 4. Cambridge University Press (1987)
3. Chien, A.Y.H., Soo, V.W.: Inferring pragmatics from dialogue contexts in simulated virtual agent games. In: *Agents for Educational Games and Simulations*, pp. 123–138. Springer (2012)
4. Fagin, R., Halpern, J.Y.: A new approach to updating beliefs. In: *Uncertainty in Artificial Intelligence*. pp. 347–374. Elsevier Science Publishers (1991)
5. Hinkelman, E.A., Allen, J.F.: Two constraints on speech act ambiguity. In: Proceedings of the 27th annual meeting on Association for Computational Linguistics. pp. 212–219. Association for Computational Linguistics (1989)
6. Núñez, R.C., Dabarera, R., Scheutz, M., Briggs, G., Bueno, O., Premaratne, K., Murthi, M.N.: DS-Based Uncertain Implication Rules for Inference and Fusion Applications. In: 16th International Conference on Information Fusion (July 2013)
7. Núñez, R.C., Scheutz, M., Premaratne, K., Murthi, M.N.: Modeling Uncertainty in First-Order Logic: A Dempster-Shafer Theoretic Approach. In: Proc. of the 8th Int. Symposium on Imprecise Probability: Theories and Applications (July 2013)
8. Pearl, J.: *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1988)
9. Perrault, C.R., Allen, J.F.: A plan-based analysis of indirect speech acts. *Computational Linguistics* 6(3-4), 167–182 (1980)
10. Pinker, S., Nowak, M.A., Lee, J.J.: The logic of indirect speech. *Proceedings of the National Academy of Sciences* 105(3), 833–838 (2008)
11. Premaratne, K., Murthi, M.N., Zhang, J., Scheutz, M., Bauer, P.H.: A Dempster-Shafer theoretic conditional approach to evidence updating for fusion of hard and soft data. In: 12th Int. Conf. on Information Fusion. pp. 2122–2129 (July 2009)
12. Searle, J.R.: Indirect speech acts. *Syntax and semantics* 3, 59–82 (1975)
13. Shafer, G.: *A Mathematical Theory of Evidence*. Princeton University Press (1976)
14. Tang, Y., Hang, C.W., Parsons, S., Singh, M.P.: Towards argumentation with symbolic Dempster-Shafer evidence. In: COMMA. pp. 462–469 (2012)
15. Wilske, S., Kruijff, G.J.: Service robots dealing with indirect speech acts. In: IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 4698–4703. IEEE (2006)