Dempster-Shafer Theoretic Learning of Indirect Speech Act Comprehension Norms

Ruchen Wen RWEN@MINES.EDU
Computer Science, Colorado School of Mines

Mohammed Aun Siddiqui AUNSID@GMAIL.COM
Thomas Williams TWILLIAMS@MINES.EDU
Computer Science, Colorado School of Mines

Abstract

For robots to successfully operate as members of human-robot teams, it is crucial that they correctly understand the intentions of their human teammates. For reasons of social courtesy, people typically employing polite utterance forms such as Indirect Speech Acts (ISAs). It is thus critical for robots to be capable of inferring the intentions behind their teammates' utterances based on both their interaction context and their knowledge of the sociocultural norms that are applicable within that context. The work described in this extended abstract, originally presented at AAAI 2020, builds off previous research on understanding and generation of ISAs using Dempster-Shafer Theoretic Uncertain Logic, by showing how Dempster-Shafer Theoretic rule learning can be used to learn appropriate uncertainty intervals for robots' representations of sociocultural politeness norms.

1. Introduction

For robots to successfully operate as members of human-robot teams, it is crucial that they correctly understand the intentions of their human teammates. This is especially important in contexts that have strong sociocultural norms, conventions, and contracts, since, in such contexts, humans typically phrase their language in the form of *Indirect Speech Acts (ISAs)* (Searle, 1975), in which the speech act's literal meaning does not match its intended meaning. For example, in a restaurant context, even though servers are in some sense functioning as subordinates to clients, it would be considered rude for restaurant-goers to say, for example, "Get me some water". Instead, restaurant-goers typically use indirect requests, such as "Could I have some water?". While this utterance may literally be a request for information, listeners effortlessly and instinctively infer the speaker's true intent, i.e., for the listener to bring them some water. Accordingly, so too must robots be able to infer the intended meanings behind their teammates' utterances according to their current context.

In previous work, we presented a Dempster-Shafer theoretic approach to ISA understanding (and generation, see Williams et al. 2015), which we argued increases the robustness of ISA inference. This approach leverages a set of Dempster-Shafer Theoretic uncertain logical rules that map utterances to inferable intentions within specified contexts. These rules are annotated with Dempster-Shafer theoretic uncertainty intervals, denoting the amount of evidence for and against (and the degree of ignorance with respect to) each rule. But while we provided mechanisms for online adaptation of these rules when corrections are explicitly provided, we did not provide any mechanisms for initially learning these intervals from observation, requiring AI practitioners to initially handcraft rules' uncertainty intervals based on their own intuitions. In this work (originally presented at the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20) (Wen et al., 2020)), we show how recent work on Dempster-Shafer Theoretic rule learning, originally developed and applied in the context of social and moral norm learning, can also be applied to learn Dempster-Shafer theoretic sociocultural politeness norms.

2. Dempster-Shafer Theortic Norm Representation and Norm Learning Problem

Dempster-Shafer theory is a belief-theoretic uncertainty-processing framework (Shafer, 1976), which has notions of belief and plausibility that are close to the inner and outer measures in probability theory (Fagin & Halpern, 2013). It is often interpreted as extending traditional Bayesian frameworks (e.g., Pearl 2014) with the ability to directly express both *uncertainty*, which is represented in the mass function, and *ignorance*, which is represented in the amount of evidence not assigned to elementary events.

We leverage DS Theory to enable uncertainty- and ignorance-sensitive representations for *pragmatic rules*: sociocultural linguistic norms usable for both ISA understanding and generation. A *Norm* $\mathcal{N}$ is an expression of the form:

$$\mathcal{N} := u \wedge C \Rightarrow i$$

where *u* represents an utterance, *C* represents a possibly empty set of contextual conditions and *i* represents a possible intention that can be inferred from utterance *u* and contextual conditions *C*.

A sociocultural linguistic *Belief-Theoretic Norm* (cf. Sarathy et al., 2017) $\mathcal{N}$ is an expression of the form:

$$\mathcal{N} := u \wedge C \Rightarrow_{[\alpha,\beta]} i$$

where each norm is associated with a Dempster-Shafer theoretic uncertainty interval $[\alpha, \beta]$ ($0 \leq \alpha \leq \beta \leq 1$).

To learn these uncertainty intervals for sociocultural linguistic norms, we use the Dempster-Shafer theoretic norm learning algorithm presented by Sarathy et al. (2017). This algorithm takes as input a finite set of data instances, and updates the beliefs and plausibilities of candidate norms as it iterates over each data instance. While Sarathy et al. (2017) originally present this algorithm in the context of learning context-sensitive social and moral norms from human data, their approach is sufficiently general that we can easily apply it to our own norm learning problem. To do so, we begin by gathering data in a two-stage experimental process similar to that used by Sarathy et al. (2017).

3. Experiments

Our first experiment was used to identify candidate ISA understanding norms. In this experiment, each participant was shown a sentence and asked to write down everything they could think of that the speaker might have meant by that sentence. 19 sentences in conventionalized ISA form were used in this experiment, chosen according to the taxonomies given by Briggs et al. (2017) and Searle (1975), and phrased to evoke a context in which speakers regularly expect sociocultural politeness norms to be used when interacting with robots (Williams et al., 2018; Foster et al., 2012). 163 U.S. participants were recruited for this experiment, yielding an average of 11.4 responses (min=7, max=17) and 4.89 unique intentions (min=3, max=7) per sentence. Across all utterances, we collected 178 total responses, comprising 25 unique intentions. Table 1 shows data collected for the sentence "*Could I have some noodles?*".

Could I have some noodles?		
Response	Logical Interpretation	Frequency
"The speaker wants the hearer to get them some noodles"	*want(S,get_for(H,S,noodles))*	5
"The speaker wants to order noodles from the hearer"	*want(S,order_from(S,H,noodles))*	4
"The speaker wants the hearer to share noodles with them"	*want(S,share_with(H,S,noodles))*	2
"The speaker wants the hearer to believe that the speaker is hungry"	*want(S,believe(H,hungry(S)))*	2

Table 1. An example of results from our first experiment.

Our second experiment collected training data that could be used to learn uncertainty intervals for the norms identified in the previous experiment. Specifically, in this experiment, we collected human judgments as to whether different intentions were appropriate for different utterances in different scenarios. These scenarios were generated based on the sentences used in the first experiment and each scenario was a combination of an environmental context and a social context. This experiment was conducted as a live, in-person laboratory study. The participant was first introduced to an experimental context, in which they were either working as a waiter, or not, and in which they were either situated in a restaurant, or at a friend's house. We presented a different experimental context in each of the two rounds for every participant. Each participant was then shown five sentences, each of which was followed by six candidate interpretations of that sentence from the first experiment. Participants were asked to select all interpretations from among those options that they believed to be interpretations of the presented sentence. Table 2 shows the six possible intentions presented for the utterance "*Can you get some noodles?*". 37 participants were recruited for our second experiment. We collected 74 data points in total, with an average of 18.5 data points (min=17, max=20) collected in each context.

4. Learning Norms from Human Data

To use the data collected in Experiment Two, we began by categorizing the data into subsets reflecting different norm frames. Specifically, four different types of norm frames were created based on the experimental contexts used in the *detection* experiment.

Type 1: Norm frames containing norms mapping utterance forms directly to intentions regardless of context. These rules' parameters can be learned from all data collected with respect to utterance u.

Candidate Intention	Logical Interpretation
The speaker wants the hearer to get them some noodles.	Int_1 = *want(S,get_for(H,S,noodles))*
The speaker wants the hearer to buy them some noodles.	Int_2 = *want(S,buy_for(H,S,noodles))*
The speaker wants to order noodles from the hearer.	Int_3 = *want(S,order_from(S,H,noodles))*
The speaker wants the hearer to believe that the speaker is hungry.	Int_4 = *want(S,believe(H,hungry(S)))*
The speaker is asking the hearer for permission whether they can have noodles.	Int_5 = *ask(S,permission(H,have(X,noodles)))*
The speaker wants the hearer to cook them some noodles.	Int_6 = *want(S,cook_for(H,S,noodles))*

Table 2. Candidate intentions presented for the utterance "*Can you get some noodles?*" in the *detection* experiment, along with the logical representations of those candidate intentions.

Type 2: Norm frames containing norms mapping utterance forms to intentions under particular environmental contexts (e.g., being in a restaurant). These rules' parameters can be learned from all data collected with respect to utterance u under that particular environmental context.

Type 3: Norm frames containing norms mapping utterance forms to intentions under particular social contexts (e.g., the listener is a waiter). These rules' parameters can be learned from all data collected with respect to utterance u under that particular social context.

Type 4: Norm frames containing norms mapping utterance forms to intentions under particular combinations of environmental and social context. These rules' parameters can be learned from all data collected with respect to utterance u under that particular combination of environmental and social context.

Under the above criteria, we gathered 5 Type 1 subsets (corresponding to the five utterances) with 74 data instances for each norm frame, 10 Type 2 subsets (corresponding to the five utterances and two environmental contexts) with 37 data instances for each norm frame, 10 Type 3 subsets (corresponding to the five utterances and two social contexts) with an average of 37 (min=35, max=39) data instances for each norm frame, and 20 Type 4 subsets (corresponding to the five utterances, two environmental contexts, and two social contexts) with an average of 18.5 (min=17, max=20) data instances for each norm frame. Note that as we progress through these types of norm frames, we attempt to learn increasingly context-specific norms from increasingly limited datasets. For each of these norm frames, we provided the data commensurate with that norm frame to Sarathy et al. (2017)'s rule learning algorithm in order to learn uncertainty intervals for each norm in that frame.

After learning uncertainty intervals for each candidate norm, we must next select a subset of "justifiable" norms that should be included in the final norm system. This can be achieved by only accepting norms that reflect a sufficient level of confidence in the consequent given the antecedent. In this paper we will examine the effects of selecting only the subset of norms for which the center of the norm's learned uncertainty interval (i.e., $\frac{\alpha+\beta}{2}$) is above some threshold τ. The left side of Figure 1 shows how the number of accepted norms declines as this threshold is reduced from 115 with threshold $\tau = 0.5$ to 6 with threshold $\tau = 0.81$. If a lower value of τ is chosen, a greater number of norms will be learned, but the agent may need to generate a greater number of clarification requests in the future; if a higher value is chosen, fewer norms will be learned but the agent may need to generate fewer clarification requests in the future. The right side of Figure 1 shows how Sarathy et al. (2017)'s norm learning algorithm converges to different uncertainty intervals for one of the norms selected with a threshold of $\tau = 0.81$.

$\mathcal{N}_1$	$QuestionYN(S, H, will(H, get(X) \wedge noodles(X))) \Rightarrow_{[0.88,0.96]} want(S, get\_for(H, noodles)$
$\mathcal{N}_2$	$QuestionYN(S, H, will(H, get(X) \wedge noodles(X))) \wedge \neg waiter(H) \Rightarrow_{[0.73,0.97]} want(S, get\_for(H, noodles)$
$\mathcal{N}_3$	$Stmt(S, H, need(H, get(S, noodles))) \wedge waiter(H) \Rightarrow_{[0.70,0.99]} want(S, get\_for(H, noodles)$
$\mathcal{N}_4$	$QuestionYN(S, H, will(H, get(X) \wedge noodles(X))) \wedge \neg in(H, restaurant) \Rightarrow_{[0.70,0.97]} want(S, get\_for(H, noodles)$
$\mathcal{N}_5$	$Stmt(S, H, should(H, get(S, noodles))) \Rightarrow_{[0.79,0.87]} want(S, get\_for(H, noodles)$
$\mathcal{N}_6$	$Stmt(S, H, need(H, get(S, noodles))) \wedge in(H, restaurant) \Rightarrow_{[0.68,0.95]} want(S, get\_for(H, noodles)$

Table 3. Six learned norms $\mathcal{N}_{1\ldots6}$ selected by our approach when a threshold of $\tau = 0.81$ was used. The convergence trajectories for $\mathcal{N}_1$ is shown below in Fig. 1.

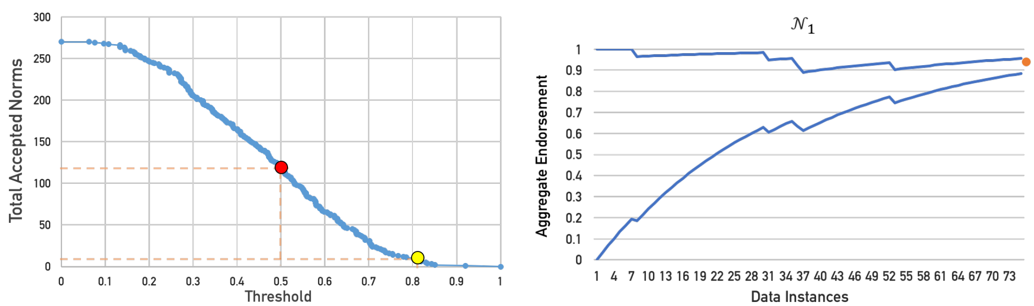

Figure 1. Left: Results from the norm learning algorithm, shows how the total of accepted norms decreases as the threshold τ increases. The red dot shows that with threshold 0.5 (selecting norms whose uncertainty intervals are centered at a point greater than 0.5) 115 total norms are selected for adoption; the yellow dot shows that when a threshold of 0.81 is used, only 6 are selected for adoption. Right: Results from the norm learning algorithm, showing convergence to different uncertainty intervals for $\mathcal{N}_1$ from Table 3. The dots represent the mean norm endorsements by experimental participants.

5. Discussion and Conclusion

We have presented the first approach to automatically learning confidence intervals for Dempster-Shafer theoretic pragmatic norms, by applying recent work on Dempster-Shafer Theoretic rule learning (Sarathy et al., 2017) to learning appropriate uncertainty intervals for robots' representations of sociocultural politeness norms. We show how this algorithm can be used with rules whose antecedents contain logically specified context descriptions of varying levels of specificity. We also demonstrated how sets of candidate norms can be selected based on the characteristics of their learned confidence intervals and how this selection process can be used to select norms that have an appropriate level of contextual specificity.

Two key challenges for future work are how to identify and generate candidate ISA understanding norms, and how to adapt those norms to facilitate lifelong learning. In this work, we selected a relatively narrow domain along with utterances that we explicitly expected to be generated due to sociocultural linguistic norms relevant to that domain. In the future it will be important to extend the methods presented in this paper to work online for long timescales over the utterances that occur naturally in human-robot dialogue, and to develop mechanisms for proposing candidate norms based on salient aspects of the robot's context. Furthermore, in future work, we plan to encode our learned norms within the pragmatic norm base (cf. Williams et al. 2015) used by the Distributed

Integrated Affect, Recognition and Cognition (DIARC) architecture (Schermerhorn et al., 2006; Scheutz et al., 2013, 2019) and assess the fluidity and task success of robots interacting with users under norm systems selected with different thresholds.

Acknowledgements

This work was supported in part by NSF grants IIS-1849348 and IIS-1909847.

References

Briggs, G., Williams, T., & Scheutz, M. (2017). Enabling robots to understand indirect speech acts in task-based interactions. *Journal of Human-Robot Interaction (JHRI)*.

Fagin, R., & Halpern, J. Y. (2013). A new approach to updating beliefs. *arXiv preprint arXiv:1304.1119*.

Foster, M. E., Gaschler, A., Giuliani, M., Isard, A., Pateraki, M., & Petrick, R. (2012). Two people walk into a bar: Dynamic multi-party social interaction with a robot agent. *Proc. ICMI*.

Pearl, J. (2014). *Probabilistic reasoning in intelligent systems: networks of plausible inference*.

Sarathy, V., Scheutz, M., Kenett, Y. N., Allaham, M., Austerweil, J. L., & Malle, B. F. (2017). Mental representations and computational modeling of context-specific human norm systems. *Proc. Annual Meeting of the Cognitive Science Society (CogSci)*.

Schermerhorn, P. W., Kramer, J. F., Middendorff, C., & Scheutz, M. (2006). DIARC: A testbed for natural human-robot interaction. *AAAI* (pp. 1972–1973).

Scheutz, M., Briggs, G., Cantrell, R., Krause, E., Williams, T., & Veale, R. (2013). Novel mechanisms for natural human-robot interactions in the DIARC architecture. *AAAI IRS Workshop*.

Scheutz, M., Williams, T., Krause, E., Oosterveld, B., Sarathy, V., & Frasca, T. (2019). An overview of the distributed integrated cognition affect and reflection DIARC architecture. In *Cog. arch*.

Searle, J. R. (1975). Indirect speech acts. *Syntax and Semantics*, *3*, 59–82.

Shafer, G. (1976). *A mathematical theory of evidence*, volume 42. Princeton university press.

Wen, R., Siddiqui, M. A., & Williams, T. (2020). Dempster-shafer theoretic learning of indirect speech act comprehension norms. *Proc. AAAI*.

Williams, T., Briggs, G., Oosterveld, B., & Scheutz, M. (2015). Going beyond command-based instructions: Extending robotic natural language interaction capabilities. *Proc. AAAI*.

Williams, T., Thames, D., Novakoff, J., & Scheutz, M. (2018). "Thank you for sharing that interesting fact!": Effects of capability and context on indirect speech act use in task-based human-robot dialogue. *Proc. 13th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*.