CS 294-115 Algorithmic Human-Robot Interaction | Fall 2016

Lecture 23: Natural Language

Scribes: Abhishek Gupta, Gregory Kahn

23.1 Language and Meaning

Natural language processing is hard because we have to figure out what words mean (e.g., “Local HS Dropouts Cut in Half”) and how words fit together (e.g., “Enraged Cow Injures Farmer With Ax”).

Languages are formalized by defining syntax and semantics. As shown in Fig. 23.1, syntax is how words are combined to form sentences and semantics is what the words mean.

Figure 23.1: Syntax and semantics

23.2 Language and Belief

Natural language processing is also hard because people don’t always say what they mean. For example, if someone asks you “Do you know what time it is?,” answering the question with “Yes” is not the desired response; what the person really wants to hear is the time. This action of implying a meaning beyond the literal sense is known as *implicature*.

Grice [1975] proposed four conversational maxims that speakers usually abide by:

Quantity: say as much as necessary, and no more

Quality: say only what is truthful

Relation: say only what is relevant

Manner: avoid unnecessary ambiguity

The question now is how to build computer systems that can draw these pragmatic inferences. We will examine two approaches: game theory and Bayesian cognitive modeling.

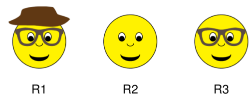

Figure 23.2: What should the speaker say in order to get the listener to select a specific face?

We will consider the example in Fig. 23.2 in which a speaker is deciding on what words to say in order to get a listener to select the correct face. The description refers to attributes (e.g., glasses), while the referent refers to the subject (e.g., R1). The listener has a probability distribution $p_L(\text{referent}|\text{description})$ and the speaker has a probability distribution $p_S(\text{description}|\text{referent})$.

Game Theory: In the game theoretic approach of "iterated best response," the listener and speaker assume the most likely response, governed by the equations in Fig. 23.4. An example run of iterated best response is also shown in Fig. 23.4. The initial belief distribution is shown in the first row. The speaker assumes the most likely response over descriptions, which corresponds to a max over the columns. The listener then assumes the most likely response over referents, which corresponds to a max over the rows. The final distribution encodes the behavior we want: referring to "hat" implies R1 while referring to glasses implies R3.

$$p_{L_i}(\text{referent}|\text{description}) = \\ \text{Unif}(\max\ p_{S_{i-1}}(\text{description}|\text{referent}))$$

$$p_{S_i}(\text{description}|\text{referent}) = \\ \text{Unif}(\max p_{L_{i-1}}(\text{referent}|\text{description}))$$

Literal listener (L0):

glasses	0.5	0	0.5
hat	1	0	0

Reasoning speaker (S1):

glasses	0	0.5	1
hat	1	0.5	0

Reasoning listener (L2):

glasses	0	0	1
hat	1	0	0

Figure 23.3: Game theory "iterated best response"

Bayes: Rational Speech Acts: In the rational speech acts theory, the speaker's goal is to convey sentences which are maximally informative (minimally surprising) to the literal listener. This can be formalized as

$$p_{L_i}(\text{referent}|\text{description}) = \\ \frac{1}{Z} p_{S_{i-1}}(\text{description}|\text{referent})$$

Given this model of the "expressive speaker," the reasoning listener uses Bayes Rule in order to infer the referent from the descriptions provided by the speaker. This corresponds to taking a norm over the rows and columns instead of the max as taken in the "iterated best response" model. An example of the rational speech acts model is shown on the following page

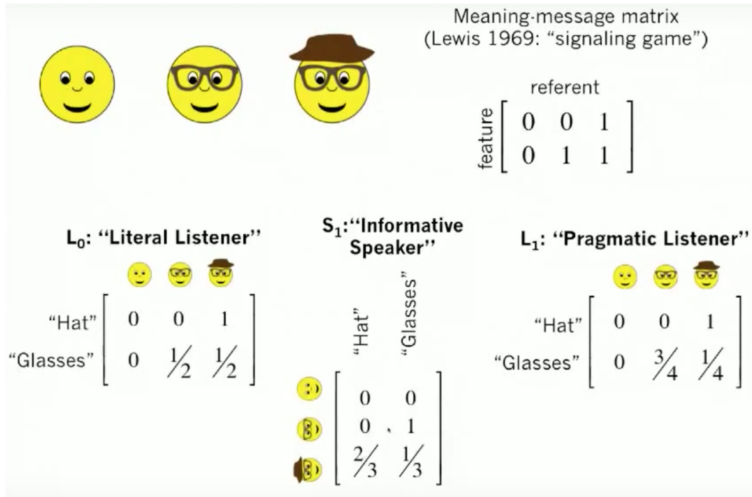

Figure 23.4: Bayes Rational Speech Acts theory depicted graphically

Further challenges in pragmatics include dealing with complex language instead of simple descriptions[Andreas] and incorporating message costs into the model.

23.3 Language and Action

Two of the major research directions in combining language and action are:

1. Question Answering (question, world) $\rightarrow$ answer
2. Instruction Following (instruction, world) $\rightarrow$ behavior

23.3.1 Behavior from language

23.3.1.1 Joint policies for reading and acting

Given a standard MDP (S,A,T,R,γ), and a set of instructions W, you can define 2^W possible attentions. This allows you to alter your MPD to be (S x 2^W, A x 2^W, T^+, R, γ). In this MDP, you can learn with any reinforcement learning or imitation algorithm. This formulation makes it easy to handle complex environments without using models, but it makes it difficult to handle hard compositional language. Good papers in this direction are [Branavan et al., 2009; Vogel Jurafsky, 2011; Mei et al., 2016].

23.3.1.2 Language to high-level plans

In this paradigm, instead of trying to learn actions directly from language, the aim is to infer descriptions of constraints and solve a planning problem. Given an initial MDP (S, A, T, R, γ), and a set of instructions W, the goal is to infer constraint structures Z and solve a planning problem such that the final plan satisfies Z intersect T. The constraint structures may look like the following

$$\lambda a.\mathrm{to}(a) \wedge \mathrm{to}(a, \iota x.\mathrm{chair}(x)) \\ \wedge \mathrm{from}(a, \iota x.\mathrm{table}(x))$$

and we are required to find plans such as "FORWARD, LEFT, FORWARD" in order to satisfy these constraints. Good papers in this direction are [Artzi et al, 2013; Tellex et al., 2011; Andreas et al 2014]

23.4 Conclusion

Pragmatics has a significant relationship with legibility as expressed in the lectures and papers on intent. Additionally language provides a naturally abstracted away hierarchical framework which we can use to inform hierarchies in our robot policies as well. Talk to Jacob Andreas for further research guidance in this direction.