Introduction to Rational Speech Act Theory

Jean Mark Gawron

May 6, 2019

Abstract

A Probabilistic Theory of Implicature Calculation

1 Introduction: Bayes' Rule

Based on the dependencies in 1, we have a joint distribution $\mathrm{P}(w,\, s \mid a)$, which can be unpacked using teh chain rule as in Sections 1.1 and 1.2.

1.1 Speaker probability

We can unpack $\mathrm{P}(w,\, s \mid a)$ using a speaker-based probability, the probability of words w given the state of the world s and the common ground a.

$$\mathrm{P}(w,\, s \mid a) \quad = \quad \mathrm{P}_{\mathrm{S}}(w \mid s,\, a) \cdot \mathrm{P}(s \mid a)$$

1.2 Listener probability

Or we can unpack $\mathrm{P}(w,\, s \mid a)$ using a listener-based probability, the probability of state s given the words w and the common ground a.

$$\mathrm{P}(w,\, s \mid a) \quad = \quad \mathrm{P}_{\mathrm{L}}(s \mid w,\, a) \cdot \mathrm{P}(w \mid a)$$

1.3 The Rule

$$\begin{array}{llll}
\text{Speaker:} & \mathrm{P}(w,\, s \mid a) & = & \mathrm{P}_{\mathrm{S}}(w \mid s,\, a) \cdot \mathrm{P}(s \mid a) \\
\text{Listener:} & \mathrm{P}(w,\, s \mid a) & = & \mathrm{P}_{\mathrm{L}}(s \mid w,\, a) \cdot \mathrm{P}(w \mid a) \\
\text{Combined} & \mathrm{P}_{\mathrm{L}}(s \mid w,\, a) \cdot \mathrm{P}(w \mid a) & = & \mathrm{P}_{\mathrm{S}}(w \mid s,\, a) \cdot \mathrm{P}(s \mid a) \\
 & \mathrm{P}_{\mathrm{L}}(s \mid w,\, a) & = & (\mathrm{P}_{\mathrm{S}}(w \mid s,\, a) \cdot \mathrm{P}(s \mid a))/\mathrm{P}(w \mid a) \\
 & \mathrm{P}_{\mathrm{L}}(s \mid w,\, a) & \propto & \mathrm{P}_{\mathrm{S}}(w \mid s,\, a) \cdot \mathrm{P}(s \mid a)
\end{array}$$

Listener Strategy: Infer the state that **maximizes** the probability of speaker's description $\mathrm{P}_{\mathrm{S}}(w \mid s,\, a)$ times the probability of what is being described $\mathrm{P}_{\mathrm{S}}(s \mid a)$.

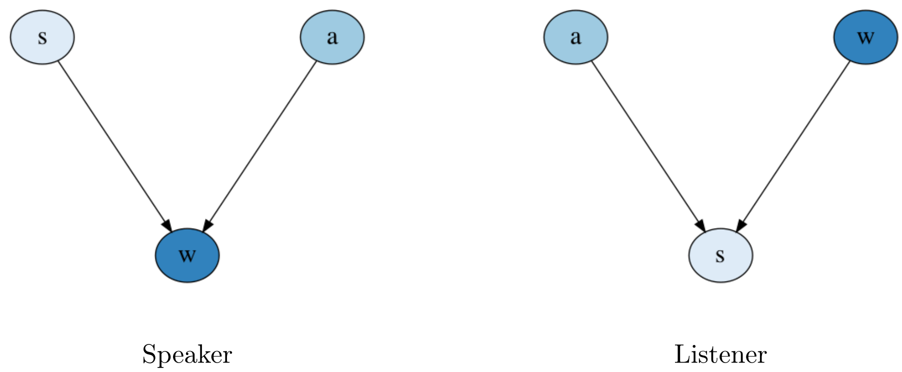

Figure 1: Communication: w = words, s = state (of the world), a = (mutually) accessible info

2 Big Picture: Baker et al. (2009), Franke et al (2016)

Baker et al. (2009) discuss a Bayesian approach to **action understanding** that is the direct forerunner of the Bayesian approach to pragmatics of Franke and Jäger (2016). We briefly describe Baker et al.

2.1 Plan understanding: Inverse planning

Plan understanding: Infer the goals of an agent based on their actions. Side effect: Will affect your probability distribution for the agent's future actions.

A simplified version of Equation (2), removing the fact that we're describing plans (sequences of states s and environments w). The left hand side describes the probability of a goal given a state and an environment:

$$\mathrm{P}(g \mid s,\, w) \propto \mathrm{P}(g \mid w)\mathrm{P}(s \mid g,\, w)$$

Bayesian rule again. Choose the g that maximizes the product on the right. We need an operationalizble way of estimating $\mathrm{P}(s \mid g,\, w)$. This comes via equation (3), soimplified again.

$$\mathrm{P}(s \mid g, w) = \sum_{a \in \mathrm{A}} \mathrm{P}(a \mid s, g, w)$$

In the experiment inspired by Gergely et al. (1995) goals are endpoints and way points in the paths taken by the sprites. Actions are paths taken to a waypoint. The model is modified to deal with sequences.

Simple v. complex goals: One condition exposed subjects to sprites that took curved paths only when there was an obstacle, and then a minimally curvbed

Table 1: Levels of analysis for pragmatic theory. Constraints are descriptions of data observations with a narrow focus. Principles are more general descriptions. Maxims are general descriptions of (speaker) behavior that aim to explain various data observations. Reasons are explanations driven by more general considerations for maxims, principles, constraints, or, directly, for data observations. Processes describe how behavior is produced: all previous levels could or should relate to processes, but do not have to do so, logically speaking.

level	example	question
constraints	Hurford's constraint: In a disjunction "*A* or *B*," *A* may not entail *B*.	what?
↕		
principles	Strongest meaning principle: Prefer strongest reading of an ambiguous sentence.	what?
↕		
maxims	Maxim of Quantity: Maximize flow of relevant information.	what/why?
↕		
reasons	Optimal language use: Be rational (or, at least, try to be)!	why?
processes	Naive serial modularity: Compute truth-conditions, then apply pragmatics.	how?

Figure 2: Levels of Explanation (Franke et al. 2016).

path. (simple) The other condition exposed subjects to sprites that naturally took curved paths even with no obstacles (complex). They were then asked to assess the probability oif test paths, one straight one curved.

2.2 Franke et al (2016)

We refer to the discussion of levels of explanation in Franke and Jäger (2016). See Figure 2.

3 Carpet Game

3.1 Preliminaries

First some facts:

Utilities	Priors	Orientation
$U_S : \begin{bmatrix} t_1^s & 0 \\ t_2^s & 20 \\ t_3^s & 40 \end{bmatrix}$	$P_L : \begin{bmatrix} t_1^s & .9 \\ t_2^s & .09 \\ t_3^s & .01 \end{bmatrix}$	$\Downarrow$ Increasingly desires to buy carpet
$U_L : \begin{bmatrix} t_1^L & 30 \\ t_2^L & 10 \end{bmatrix}$	$P_S : \begin{bmatrix} t_1^L & .9 \\ t_2^L & .1 \end{bmatrix}$	$\Downarrow$ Increasingly eager to sell carpet

Each utility function give us a value for the carpet, which we call V_S for S [the buyer] and V_L for L [the seller]. They represent different facts. V_S is the highest price S is willing to pay, and it varies with her state of mind. V_L is the lowest price L will accept, and it also varies. Each Prior represents a player's view of the **opposing** player's mind set. Thus $P_L$is L (the seller)'s estimate of how eager S is to buy the carpet, which is rather pessimistic, so as to make our story go.

Expected values for V_S and V_L, based on the player's priors.

$$\begin{array}{rcccccc} E(V_S) & = & P_L(t_1^s) \cdot U_S(t_1^s) & + & P_L(t_2^s) \cdot U_S(t_2^s) & + & P_L(t_3^s) \cdot U_S(t_3^s) \\ & = & .9 \cdot 0 & + & .09 \cdot 20 & + & .01 \cdot 40 \\ & = & & & 2.2 & & \\ E(V_L) & = & P_S(t_1^L) \cdot U_L(t_1^L) & + & P_S(t_2^L) \cdot U_L(t_2^L) & & \\ & = & .9 \cdot 30 & + & .1 \cdot 10 & & \\ & = & & 28 & & & \end{array}$$

Note: $E(V_L) \not\leq E(V_S)$, so we're not conforming to the basic picture in which there is a price on which S and L can agree:

$$E(V_L)\_\_\_\_\_\text{Price}\_\_\_\_E(V_S)$$

3.2 Why Expected Values?

Let's assume for the sake of argument that there is a fact of the matter, a lowest price L will agree to and a highest price S will pay. Why do we compute **the expected values of** $E(V_L)$ and $E(V_S)$ rather than just using those actual values? Our assumption is that L does not know V_S and S does not know V_L. We use the expected values to reason about their best strategies given this state of uncertainty. Moreover, even though L **does** know V_L, she uses $E(V_L)$ to reason about S's best strategies.

Here's an example. Based on this belief state for S and L, L's best strategy in F&J's Figure 5 is to say nothing. Why? Because whether L offers the high price or the low price, S's best option is to reject. S's payoff for an ACCEPT on the Low Branch is -13.8; S's payoff for an ACCEPT on the High Branch is -33.8. In both cases S's payoff for a REJECT is better (-1).

3.3 Updating the prior

But here is the thing. We have the power with language to change other people's belief states. The sentences in F&J's (9), repeated here, represent S's options for trying to change P_L, L's prior distribution over S's degree of interest. Note that s_1 and s_2 are somewhat indirect as regards the question of whether S wants to buy, while s_3 is direct.

(9) s_1 No This rug has somewhat faded colors, but the pattern is kind of nice.
s_2 No This is a beautiful carpet.
s_3 Yes I have decided to buy this carpet.

Note that they represent an increasing degree of interest in buying the carpet. We'll assume that L is naive, and does not consider the possibility of being manipulated (somewhat against the stereotype for carpet retailers),

F&J's Table 3, reproduced here. represents some made-up numbers for how the utterance options (s_1, s_2, and s_3) change P_L. The magnitudes are contrived so as to make something happen, but the relaive sizes encode simple intuitions:

Expressing increasing
S desire for carpet
$\Rightarrow$

$P_S(t^S_j \mid s_i)$:

	s_1	s_2	s_3	
t^S_1	$9x$	x	0	
t^S_2	$4x$	$30x$	$20s$	$\Downarrow$ Actual S desire
t^S_3	x	$150x$	$500x$	

The idea here is that an utterance of one of these three sentences by S can influence L's prior distribution over S's state, P_L, and that in turn can change all the

3.4 Bayes' Rule Explains F&J's Table 4

Table 4 is gotten from table 3 by applying Bayes' Rule. The version of Bayes' Rule we'll need is the same as the version above, only variables have to be relabeled. Here's the relabeled version:

$$P_L(t^S_1 \mid s_1) \propto P_S(s_1 \mid t^S_i) \cdot P_L(t^S_i)$$

		$P_S(s_1 \mid t_i^s)$	$\cdot$	$P_L(t_i^s)$		$P_L(t_i^s \mid s_1)$
$P_L(t_1^s \mid s_1)$	$\propto$	$9x$	$\cdot$	.90	$8.1x$	.956
$P_L(t_2^s \mid s_1)$	$\propto$	$4x$	$\cdot$	.09	$.36x$	.042
$P_L(t_3^s \mid s_1)$	$\propto$	$1x$	$\cdot$	.01	$.01x$	.001
		Total			$8.47x$	~ 1.000

Bayesian update: $P(t_i^s)$ before and after s_i

	Prior	Posteriors		
	$P(t_i^s)$	$P(t_i^s \mid s_1)$	$P(t_2^s \mid s_2)$	$P(t_3^s \mid s_3)$
t_1^s	.90	.956	.18	.00
t_2^s	.09	.043	.53	.29
t_3^s	.01	.001	.26	.74
	p. 31	Table 4, p. 33		

Of course, updating a prior will have an immediat effect of V_S. Let's assume S says s_2 and recompute V_S

$$
\begin{aligned}
E(V_S) &= P_L(t_1^s \mid s_1) \cdot U_S(t_1^s) + P_L(t_2^s \mid s_1) \cdot U_S(t_2^s) + P_L(t_3^s \mid s_1) \cdot U_S(t_3^s) \\
&= .18 \cdot 0 + .53 \cdot 20 + .26 \cdot 40 \\
&= 22.2
\end{aligned}
$$

The expected value of V_S was 2.2, based on P_L, which represents L as having a rather pessimistic view of V_S. Because of s_2, it has now risen to 22.2. According to the game analysis in Figure 4, S's payoff for an ACCEPT in the low branch now becomes positive. In fact, numerous changes occur on both the low and high branches with each s_i. The following table summarizes the payoff situation (see also F&J's Figure 6):

Expected S Payoffs: Offer $- V_S$

	s_1		s_2		s_3	
Offer	V_S	Payoff	V_S	Payoff	V_S	Payoff
15	0.8	−14.2	22.2	7.2	34.8	19.8
35	0.8	−34.2	22.2	−12.8	34.8	−.2

If L offers the low price $15, S will accept after s_2, because the payoff of 7.2 is better than the payoff of -1 with REJECT. Of course whether L will do this depends on the **actual** V_L, because L's payoff on the low price is $-V_L + 15$, so he's only going to name the low price if he's in state t_2^L, where he's quite eager to sell. F&J's Figure 7 lays out the whole picture for L's best strategy, given each of the possible utterances and each possible state of L.

F&J's Figure 7 makes it clear that S's expectations of success depend on L's state. We can compute expected values for S's strategies using P_S

$$P_S: \begin{bmatrix} t_1^L & .9 \\ t_2^L & .1 \end{bmatrix}$$

and the best response analysis in F&J's Figure 7, which tell us what a naive rational L will do in each of his 2 possible states.

Figure 8 provides S's expected payoff for the 3 possible values V_S. Let's take as an example the third branch, t_3^S, in which S wants the carpet the most, $V_S = 40$. Each of L's choices pays off S based on the rule given in F&J's Figure 4:

L's Strategy	Rule	In t_3^S
Low	$V_S - 15$	25
High	$V_S - 35$	5
No offer	0	0

With each possible utterance s_i and each t_i^L, the optimal L strategies come from F&J's Fig 7.

s_1	$P_S(t_1^S) \cdot$ [no offer]	$+$	$P_S(t_2^S) \cdot$ [no offer]
	$.9 \cdot 0$	$+$	$.1 \cdot 0$
	0		
s_2	$P_S(t_1^S) \cdot$ [no offer]	$+$	$P_S(t_2^S) \cdot$ [Low]
	$.9 \cdot 0$	$+$	$.1 \cdot 25$
	2.5		
s_3	$P_S(t_1^S) \cdot$ [High]	$+$	$P_S(t_2^S) \cdot$ [High]
	$.9 \cdot 5$	$+$	$.1 \cdot 5$
	5		

On this branch clearly s_3 clearly yields the highest payoff. Any other utterance runs too much risk of a no offer.

The other branches in Figure 8 are computed in a similar way. What should S say? It clearly depends on S's state. Each branch of Figure 8 gives a different answer.

t_i^S	V_S	Best strategy
t_1^S	0	s_1
t_2^S	20	s_2
t_3^S	40	s_3

Now s_1, s_2, s_2 So the more S values the carpet, the more desire she should express for it.

This story has several lessons:

Figure 1: Example context for a reference game trial, after (Frank and Goodman 2012).

Figure 3: Reference Games

1. Variability of what a message signals
2. Role of messages in creating more favorable posterior
3. Variability of what a speaker should choose to say based on speaker's state t_i^{s}, and their assessment of listener's state ($\mathrm{P_s}$).
4. Indirectness explained: They discuss an alternative starting prior they call the **eagerness assumption**, on which S assumes L is much more eager to sell (F&J's Figure 9), on which it pays for even a strongly interested buyer to be indirect (use s_2 in favor of the direct s_3).

4 Reference Games

$\mathrm{P_{literal}}$:

	property			
referent	square	green	circle	blue
Gr. Sq.	1.0	.5	0	0
Gr. Circ.	0	.5	.5	0
Bl. Circl	0	0	.5	1

Read this as

$$\mathrm{P_{literal}}(\text{Gr. Sq.} \mid \text{green}) = .5$$

In general they write $\mathrm{P_{literal}}(r \mid p)$, r for referent, p for property.

Add in a speaker preference function f:

$$\mathrm{EU_{speaker}}(\text{refer to r, choose p, paraneter f}) = \mathrm{P_{literal}}(r \mid p) + f(p)$$

This number is no longer a probability.

Define a probability function **based on $\mathbf{EU_{speaker}}$**) using **softmax**:

$$\mathrm{P_{prod}}(p \mid r; \lambda, f) = \frac{\exp(\lambda \cdot \mathrm{EU_{speaker}}(r, p; f))}{\sum_{p'} \exp(\lambda \cdot \mathrm{EU_{speaker}}(r, p'; f))}$$

Softmax idea:

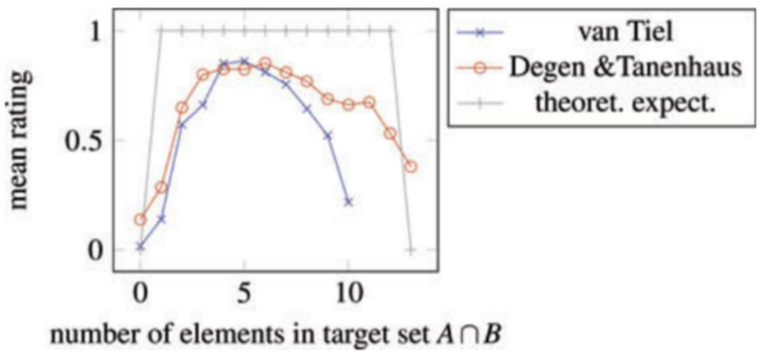

Figure 3: Mean Likert-scale ratings of sentences of the form "Some of the *A*s are *B*s" for different target set sizes in different experiments (see main text). The gray line gives the expected applicability under standard linguistic theory (for $|A| = 14$).

Table 1: Scalar some

1. A probability function guided by a utility function
2. The optimal choice is the most probable
3. Suboptimal choices are possible. A suboptimal choice c is more probable than any choice with less utility.

Bottom line:

1. Fix parameters λ and x as the choice that makes the observed dat a most probable
2. Different parameters for productiona dn comprehensiopn experiments
3. Great fit
4. Toy example demonstrating only that probabilistic pragmatics provides an empirically testable framework modeling production and comprehnsion choices

5 Gradience

1. Scalar implicatures are variable:Differences between *some/all* and *big/enormous* or *attractive/stunning*
2. **Patterns of variation**

- Given 10 circles, told *circles are white*, asked how many?
- 4 or 5 very common, 2 less likely than 6.
- Likert Scale judgments of how well sentences describe a situation. Vary $| A \cap B |$.
- The notion prototype seems more helpful.

6 Wonky Worlds

Degen et al. (2015) **?** deals with two important predictions of the RSA model.

(1) a. Some of the marbles sank.

b. All the marbles sank.

c. Some of the cars sank.

d. All of the cars sank.

1. What is Θ_X? Or Θ_{marble} in this case?
2. Suppose Θ_X is not extreme, is P(1b) high or low after an utterance of (1a)?
3. As Θ_X approaches 1, what happens to P(1b) [according to the RSA model]?
4. As Θ_X approaches 1, what happens to P(1b) [in experiments with human subjects]?
5. For moderate to high Θ_{car} and a high number of cars, is P(1d) substantially different after an utterance of P(1c) [according to the RSA model]? [according to human subjects]?
6. What does the wonky worlds parameter do?
7. Moral of these experiments: What should researchers think about when it comes to using "odd" items in semantics and pargmatics experiments?

7 Potts et al. (2016)

7.1 Exhaustification

Exhaustification of ϕ is defined versus some set of alternatives. Exhaust(ϕ) is ϕ conjoined with the negation of all the alternatives that are stronger tha ϕ.

For example the exhaustification of $\exists x\, \phi(x)$ in a context where $\forall x\, \phi(x)$ is the only alternative will be equivalent to

$$\exists x\, \phi(x)\, \&\, \sim \forall x\, \phi(x)$$

because

$$\forall x\, \phi(x) \Rightarrow \exists x\, \phi(x)$$

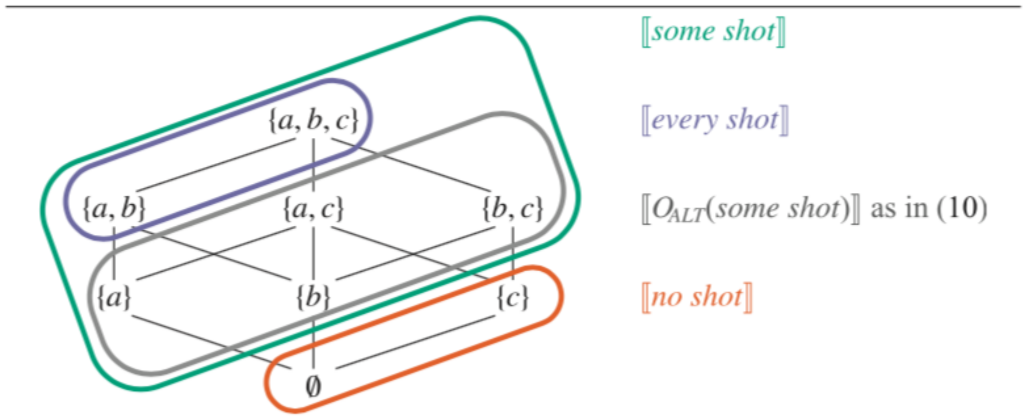

Figure 1 Given a domain {*a*,*b*,*c*} with ⟦*shot*⟧ = {*a*,*b*}, ⟦*some shot*⟧ is equal to the set of sets in the green box, ⟦*every shot*⟧ to the set of sets in the purple box, and ⟦*no shot*⟧ to the set of sets in the orange box. If *ALT*(*some shot*) contains ⟦*every shot*⟧, then *some shot* is refined to exclude the purple subset.

Figure 4: Exhaustification (Potts et al 2016)

References

Baker, Chris L, Rebecca Saxe, and Joshua B Tenenbaum. 2009. Action understanding as inverse planning. *Cognition* 113(3):329–349.

Franke, Michael, and Gerhard Jäger. 2016. Probabilistic pragmatics, or why bayes' rule is probably important for pragmatics. *Zeitschrift für sprachwissenschaft* 35(1):3–44.

Gergely, György, Zoltán Nádasdy, Gergely Csibra, and Szilvia Bíró. 1995. Taking the intentional stance at 12 months of age. *Cognition* 56(2):165–193.

Potts, Christopher, Daniel Lassiter, Roger Levy, and Michael C Frank. 2016. Embedded implicatures as pragmatic inferences under compositional lexical uncertainty. *Journal of Semantics* 33(4):755–802.