

An Incremental Iterated Response Model of Pragmatics

Reuben Cohn-Gordon

Stanford

Noah D. Goodman

Stanford

Christopher Potts

Stanford

{reubencg, ngoodman, cgpotts}@stanford.edu

Abstract

Recent *Iterated Response* (IR) models of pragmatics conceptualize language use as a recursive process in which agents reason about each other to increase communicative efficiency. These models are generally defined over complete utterances. However, there is substantial evidence that pragmatic reasoning takes place incrementally during production and comprehension. We address this with an incremental IR model. We compare the incremental and global versions using computational simulations, and we assess the incremental model against existing experimental data and in the TUNA corpus for referring expression generation, showing that the model can capture phenomena out of reach of global versions.

1 Introduction

A number of recent Bayesian models of pragmatics conceptualize language use as a recursive process in which abstract speaker and listener agents reason about each other to increase communicative efficiency and enrich the meanings of the utterances they hear in context-dependent ways (Jäger 2007, 2012; Franke 2009; Frank and Goodman 2012; for overviews, see Franke and Jäger 2014; Goodman and Frank 2016). For example, in these models, pragmatic listeners reason, not about the literal semantics of the utterances they hear, but rather about pragmatic speakers reasoning about simpler listeners that are defined directly in terms of the literal semantics. In this back-and-forth, many phenomena characterized by Grice (1975) as *conversational implicatures* emerge naturally as probabilistic inferences.

In general, these *iterated response* (IR) models separate pragmatic reasoning from incremental processing, in that the calculations are done in terms of complete utterances. However, there is substantial evidence that pragmatic processing

is incremental: listeners venture pragmatic inferences over the time-course of the utterances they hear, which influences the choices that speakers make. To address this, we develop an IR model that is incremental in the sense that pragmatic reasoning takes place word-by-word (though the process could be defined in terms of different linguistic units, like morphemes or phrases).

A variant of this model was applied successfully to *pragmatic image captioning* by Cohn-Gordon et al. (2018); here we concentrate on its qualitative behavior and linguistic predictions. We present computational experiments which demonstrate that incremental and global pragmatics make different predictions, and we show that a speaker that incrementally makes pragmatically informative choices arrives at an utterance which is globally informative. We then argue that an incremental model can account for two empirical observations out of reach of a global model: (i) the asymmetry between adjective–noun and noun–adjective languages in over-informative referential behavior (Rubio-Fernández, 2016), and (ii) the anticipatory implicatures arising from contrastive modifiers (Sedivy, 2007). The first of these observations requires a model of language production, while the second requires a model of language interpretation, and as such these case studies serve to demonstrate both aspects of incremental pragmatics. Finally, we apply the model to the TUNA corpus for referring expression generation (Gatt et al., 2009), showing that it makes more realistic predictions about attributive modifiers than does its global counterpart.

2 Iterated Response Models

We construct our model within the *Rational Speech Acts* (RSA) paradigm (Frank and Goodman, 2012; Goodman and Stuhlmüller, 2013).

RSA and its extensions have been applied to a wide range of pragmatic phenomena, including scalar implicatures (Frank et al., 2016; Potts et al., 2016), manner implicatures (Bergen et al., 2016), hyperbole (Kao et al., 2014), metaphor, and politeness (Yoon et al., 2016). In addition, RSA can be cast as a machine learning model, thereby allowing us to study pragmatic reasoning in large corpora and complex environments (Vogel et al., 2013; Monroe and Potts, 2015; Monroe et al., 2017; Andreas and Klein, 2016).

Standard RSA models are global in the sense that the pragmatic reasoning is defined over complete utterances. Speakers are conditional distributions of the form $P(u|w)$, while listeners are of the form $P(w|u)$, for an utterance u and state w . We first present this global formulation (section 2.1), and then we show how to reformulate it so that utterances are sequences of linguistic units $u = [u_1, \dots, u_n]$ and the core RSA reasoning is applied to each step u_i given $[u_1, \dots, u_{i-1}]$.

Figure 1 presents a running illustrative example. We imagine there are three referents, a red dress (R1), a blue dress (R2), and a red hat (R3). We have a simple language composed of three utterances, *dress*, *red dress*, and *red object*, each with its expected semantics. For the RSA calculation, we make the background assumption that the speaker and listener are playing a coordination game: they succeed to the extent that the listener can use the speaker’s utterance to identify the speaker’s intended referent.

2.1 Global Pragmatics

We define our global RSA agents as follows:

$$L_0^{\text{UTT}}(w|u) \propto \llbracket u \rrbracket(w) \quad (1)$$

$$S_1^{\text{UTT-GP}}(u|w) \propto e^{\log(L_0^{\text{UTT}}(w|u)) - \text{cost}(u)} \quad (2)$$

$$L_1^{\text{UTT}}(w|u) \propto S_1^{\text{UTT-GP}}(u|w) \quad (3)$$

Here, $\llbracket \cdot \rrbracket$ is an interpretation function mapping utterances to functions from referents to $\{0, 1\}$. Thus, the literal listener L_0^{UTT} is simply a probabilistic version of the truth conditions established by $\llbracket \cdot \rrbracket$; given an utterance u , L_0^{UTT} evenly distributes probability mass to the worlds compatible with u according to $\llbracket \cdot \rrbracket$.

The pragmatic speaker $S_1^{\text{UTT-GP}}$ is more sophisticated than a literal agent, in that $S_1^{\text{UTT-GP}}$ reasons about L_0^{UTT} , taking message costs into account. We take *cost* to be a language model, which could either be estimated from data (higher probability

to attested utterances), derived from a grammar (higher probability to grammatical utterances), or simply assign longer utterances more cost. Intuitively, $S_1^{\text{UTT-GP}}$ prefers utterances which are not only true but best convey to L_0^{UTT} which world the speaker is in. This is illustrated in figure 1b: whereas all three messages are true of R1, $S_1^{\text{UTT-GP}}$ prefers *red dress* because it is the most specific. In this sense, $S_1^{\text{UTT-GP}}$ is a model of a Gricean informative speaker.

The pragmatic listener L_1^{UTT} in turn reasons about what world state $S_1^{\text{UTT-GP}}$ must be in such that the observed utterance was chosen, and thus draws more refined inferences than L_0^{UTT} . We see this in figure 1b as well, with respect to *dress* and *red object*. Whereas L_0^{UTT} regards these messages as completely ambiguous, L_1^{UTT} (softly) disambiguates them: *dress* is heavily biased toward R2, and *red object* is heavily biased toward R3. This inference formalizes the intuitive reasoning that if the speaker of *red object* had been referring to R1, they would have used the more specific, informative *red dress*; their avoidance of this means they must be referring to R3.

2.2 Incremental Pragmatics

In natural language, speakers and listeners produce and comprehend utterances segment by segment. For present purposes, we define this process at the word level, but we emphasize that the proposed approach extends both to sub-word segments (Cohn-Gordon et al., 2018) and to larger syntactic units.

To approximate this incremental process, we represent utterances as sequences of words and allow RSA-style reasoning to happen at the point of production or comprehension of each word, in the order they are uttered. We represent the end of an utterance as a STOP token, so that the choice of STOP as the next “word” represents the decision that the utterance is complete.

Roughly speaking, we want to define listener models $P(w|word, c)$ and speaker models $P(word|w, c)$, where c is a sequence of words constituting the utterance so far. In order to do this, we first must define an incremental semantics. This incremental semantics is defined in terms of a global semantics and the set of available complete utterances. For any partial sequence c and set of referents W , $\llbracket c \rrbracket(w) \in [0, 1]$ is the number of full-utterance extensions of c true in w divided




$\llbracket \cdot \rrbracket$	R1	R2	R3	<i>cost</i>
<i>dress</i>	1	1	0	<i>dress</i> 0
<i>red dress</i>	1	0	0	<i>red dress</i> 0
<i>red object</i>	1	0	1	<i>red object</i> 0


(a) Reference game.

L_0^{UTT}	R1	R2	R3	$S_1^{\text{UTT-GP}}$	<i>dress</i>	<i>red dress</i>	<i>red object</i>	L_1^{UTT}	R1	R2	R3
<i>dress</i>	0.5	0.5	0.0	R1	0.25	0.5	0.25	<i>dress</i>	0.2	0.8	0
<i>red dress</i>	1.0	0.0	0.0	R2	1.0	0.0	0.0	<i>red dress</i>	1.0	0.0	0.0
<i>red object</i>	0.5	0.0	0.5	R3	0.0	0.0	1.0	<i>red object</i>	0.2	0.0	0.8

(b) Global RSA.



(c) Incremental RSA speaker predictions.

(d) Incremental RSA listener predictions upon hearing *red*.

$S_1^{\text{UTT-IP}}$	<i>dress</i>	<i>red dress</i>	<i>red object</i>
R1	0.42	0.38	0.20
R2	1.0	0.0	0.0
R3	0.0	0.0	1.0

(e) Incremental utterance-level predictions from $S_1^{\text{UTT-IP}}$.

Figure 1: Illustrative example comparing global and incremental RSA. For ease of comparison to the global model, we do not depict the STOP token for the incremental model.

by the number of possible extensions of c into full utterances that are true of any world in W . Where c is a full utterance, $\llbracket c \rrbracket(w) \in \{0, 1\}$ is as in the global model; where c is a partial utterances, $\llbracket c \rrbracket$ represents the biases created by c .

This is not the only possible way to define an incremental semantics. One alternative is to define a probabilistic S_0 or L_0 directly (as in Cohn-Gordon et al. 2018). We choose the method defined above since it is as close as possible to a standard truth-conditional semantics and thus permits direct comparison. Furthermore, since our incremental semantics can be generated from an utterance-level semantics, only the latter needs to be stipulated.

Just as the standard RSA model presented in section 2.1 uses as global semantics to define successive speakers and listeners, we can now define an incremental literal listener L_0^{WORD} and incre-

mental pragmatic speaker and listener S_1^{WORD} and L_1^{WORD} :

$$L_0^{\text{WORD}}(w|c, \text{word}) \propto \llbracket c + \text{word} \rrbracket(w) \quad (4)$$

$$S_1^{\text{WORD}}(\text{word}|c, w) \propto e^{\log(L_0^{\text{WORD}}(w|c, \text{word})) - \text{cost}(\text{word})} \quad (5)$$

$$L_1^{\text{WORD}}(w|c, \text{word}) \propto S_1^{\text{WORD}}(\text{word}|c, w) \quad (6)$$

Figure 1c summarizes the reasoning of the incremental pragmatic speaker S_1^{WORD} , assuming 0 cost on all words for simplicity. This agent prefers *red* as a first word when conveying R1: $S_1^{\text{WORD}}(\text{red}|c = \llbracket \cdot \rrbracket, w = \text{R1}) = 0.57$. However, if R3 is the intended referent, the agent *must* begin its utterance with *red* (since *hat* is not an available word in this simple example). As shown in figure 1d, this fact allows the pragmatic listener to infer from hearing *red* that the referent is most likely R3: $L_1^{\text{WORD}}(\text{R3}|c = \llbracket \cdot \rrbracket, \text{red}) = 0.64$.

There may be cases in which there is no possible true continuation of a sequence of words into a true utterance. For instance, no continuation of *red* constitutes a truthful description of R2. In such situations, we say that probability is evenly distributed over all choices of word, so that $S_1^{\text{WORD}}(\text{dress}|c = [\text{red}], w = R2) = S_1^{\text{WORD}}(\text{object}|c = [\text{red}], w = R2) = 0.5$.

2.3 An Utterance-level Incremental Speaker

From the word level agent S_1^{WORD} , we can use the chain rule to obtain $S_1^{\text{UTT-IP}}$, an *utterance-level speaker* whose values are the result of incremental pragmatic inferences:¹

$$S_1^{\text{UTT-IP}}(u|w) = \prod_{i=1}^n S_1^{\text{WORD}}(u_i|c = [u_1 \dots u_{i-1}], w) \quad (7)$$

Whereas $S_1^{\text{UTT-GP}}$ in (2) makes pragmatic calculations on the basis of whole utterances, $S_1^{\text{UTT-IP}}$ makes incremental pragmatic decisions about each choice of word, which together also give rise to a distribution over utterances.

This allows for an efficient strategy, namely *greedy unrolling*, to generate an utterance from a referent r . We choose the first word $word_1$ of the utterance to be $\arg \max_{word} S_1^{\text{WORD}}(word|w = r, c = [])$, and this decision then becomes part of the context for choosing the second word: $word_2$ is $\arg \max_{word} S_1^{\text{WORD}}(word|w = r, c = [word_1])$. And so on through the entire utterance.

2.4 Relating the Global and Incremental Models

Figure 2 depicts the core relationships between the global and local models, focusing on the pragmatic speaker. The agents along the solid green path define the global model of section 2.1, while those along the dashed red path define the incremental model of section 2.2 as defined by $S_1^{\text{UTT-IP}}$. Importantly, while $S_1^{\text{UTT-GP}}$ and $S_1^{\text{UTT-IP}}$ are of the same *type*, in the sense of being conditional probability distributions over full utterances, they are not the same distribution.

For instance, the predictions of $S_1^{\text{UTT-IP}}$ for our illustrative example are given in figure 1e. Comparing them with the global pragmatic speaker predictions in figure 1b, we see that the two make

¹We use $u[n]$ for the n th element of a list u , and $u[:n]$ for the sublist of u up to but not including $u[n]$.




Figure 2: Two ways of constructing an utterance-level pragmatic speaker from a semantics. The solid green path is to obtain a literal listener over full utterances and then perform pragmatics, which gives rise to $S_1^{\text{UTT-GP}}$ while the dashed red path is to obtain an incremental literal listener, use it to construct a word-level pragmatic speaker from L_0^{WORD} and then use this to define an utterance-level pragmatic speaker, $S_1^{\text{UTT-IP}}$.

substantively different predictions. In the global model, the speaker who wishes to refer to R1 prefers *red dress*. In contrast, in the incremental model, the speaker referring to R1 prefers *dress*. The reason for $S_1^{\text{UTT-IP}}$ having these values is that saying *dress* ensures the termination of the utterance (given the set of utterances that are available in this example), which therefore has probability 1.0 at the next time step, while saying *red* leaves two options, *dress* and *object*. Thus, $S_1^{\text{UTT-GP}}$ and $S_1^{\text{UTT-IP}}$ are not only quantitatively different, but even differ in their predictions about which utterances are optimal.

Figure 3 provides an abstract example which further reinforces this difference. Moving from left to right, the numbers in green depict the S_1^{WORD} probabilities at each of the two steps in the generation of a complete utterance, when the target reference is W1 instead of distractor W2. The probability of the full utterance at $S_1^{\text{UTT-IP}}$ is the product of the two S_1^{WORD} steps.

An example of the difference to $S_1^{\text{UTT-GP}}$ is shown in green. When referring to W1, $S_1^{\text{UTT-GP}}$ gives equal weight to AA, BA and BB. $S_1^{\text{UTT-IP}}$, however, first chooses between A and B: in this decision, B is preferred, since one of the two continuations of A, namely AB, is not compatible with W1. However, if A is chosen, the subsequent choice is fully determined to be B ($p(A|[A], W1 = 1.0$). This results in a preference for AA.

While we focus largely on the differences between incremental and global pragmatics, it is worth highlighting a regard in which the former behaves like the latter.

Given a referent r , call an utterance u *weakly*




Figure 3: A depiction of the probabilities of the $S_1^{\text{UTT-IP}}$ in red and $S_1^{\text{UTT-GP}}$ in green, for a simple abstract example. The four utterances are AA , AB , BA and BB , while the two worlds are $W1$ and $W2$. The semantics assigns $u = W1$ and $w = AB$ to 0 but all other utterance–world pairs to 1. Given a world w , the incremental speaker first chooses the first letter to be A or B , and then chooses the second letter conditioned jointly on w and the first letter, to obtain a full utterance. The resulting full utterance probabilities are compared with the predictions of the global $S_1^{\text{UTT-GP}}$. As can be seen, the incremental and global speakers assign different probabilities to each utterance and are consequently distinct from each other.

informative if $L_0^{\text{UTT}}(r|u) \geq \frac{1}{|W|}$, where W is the set of possible referents. In other words, given u , the literal listener L_0^{UTT} will guess the correct referent with probability at least at chance (when costs are 0). We note that the utterance $u*_r$ obtained by greedy unrolling at each step of generation, as described in section 2.3, is weakly informative. To see this, observe that the n th word of $u*_r$ is $\arg \max_{\text{word}} S_1^{\text{WORD}}(\text{word}|w = r, c = u*_r[:n])$. Since at each step S_1^{WORD} produces a word which, at worst, does not rule out any referents for L_0^{WORD} , the resulting sentence $u*$ at worst gives $L_0^{\text{UTT}}(r|u*_r) \geq \frac{1}{|W|}$. In other words, greedily unrolling the incremental speaker will produce an utterance which is true and has the literal listener infer the probability of the intended referent as being at least at chance.

This result suggests that the strategy of choosing the most informative word (or syntactic unit) at each point in the generation of an utterance can be used as a substitute for choosing, from all utterances, the one which is most informative. Exploring the gap in this informativity bound is an important future direction.

One potential advantage of $S_1^{\text{UTT-IP}}$ as a plausible model of referring expression generation (over a small, discrete set of referents) is computational

tractability in real-world settings where the set of possible utterances U is large or unbounded. In such settings, $S_1^{\text{UTT-GP}}$ becomes intractable, owing to the normalizing term required over U .

However, it is important to note that in the setup presented here, $S_1^{\text{UTT-IP}}$ is not more tractable than $S_1^{\text{UTT-GP}}$. The reason for this is that the calculation of the incremental semantics depends on factoring the global semantics: this involves an intractable normalization akin to that present in the $S_1^{\text{UTT-GP}}$. Thus, $S_1^{\text{UTT-IP}}$ is only computationally tractable if the incremental semantics in terms of which it is defined is tractable. In the present work, we define an incremental semantics in terms of a global one, in order to allow for a maximally clear comparison between the global and incremental models of pragmatics that ensue. However, an incremental semantics (or literal speaker model) can be defined or learned independently, as in [Vedantam et al. 2017](#) and [Cohn-Gordon et al. 2018](#).

3 Application to Prior Experiments

We now briefly consider two cases where incremental pragmatics provides an explanation of a phenomenon where global pragmatics does not seem to suffice.

3.1 Over-informative Referring Expressions

It has been observed that, when generating referring expressions (REs), humans often provide more information than necessary to refer unambiguously (Engelhardt et al., 2006; Hermann and Deutsch, 1976). For instance, Rubio-Fernández (2016) shows that English speakers often use redundant color terms (e.g., *the red dress*) in a scene with only a single dress, where the shorter utterance *dress* would suffice. However, Rubio-Fernández also notes that Spanish speakers are less likely to over-describe with the analogous referring expression, *el vestido rojo*, in the same situation. This difference is a challenge for non-incremental pragmatic accounts, since, *ceteris paribus*, we would expect semantically equivalent Spanish and English REs to have the same production probability.

Using incremental pragmatics, we model the English case as follows: let the referents be a red dress (R1) and a blue hat (R2), and the possible utterances be *dress*, *red dress*, *hat*, and *blue hat*, with the obvious semantics.

We make the following assumption regarding the *cost* term: assume a cost of 1.0 for all words but a cost of 0.0 for the STOP token. Further assume that an utterance’s cost is the sum of the cost of its words. The effect of this cost term is to penalize longer utterances, all else being equal.

On these assumptions, the globally pragmatic speaker $S_1^{\text{UTT-GP}}$ prefers *dress* to *red dress*, since both are fully informative but the latter is costlier: $S_1^{\text{UTT-GP}}(\textit{dress}|\text{R1}) = 0.73 > S_1^{\text{UTT-GP}}(\textit{red dress}|\text{R1}) = 0.27$. Meanwhile the incremental pragmatic speaker $S_1^{\text{UTT-IP}}$ is undecided: $S_1^{\text{UTT-IP}}(\textit{dress}|\text{R1}) = S_1^{\text{UTT-IP}}(\textit{red dress}|\text{R1}) = 0.5$. The increase in mass on the over-informative RE *red dress* in $S_1^{\text{UTT-IP}}$ as compared to $S_1^{\text{UTT-GP}}$ is the result of incremental processing: the decision between *red* and *dress* is made on the basis of informativity, and both words are equally informative. However, if *red* is chosen, the subsequent, now over-informative word *dress* has to follow, since *red* on its own is not an utterance.

We explore the generality of this dynamic – that incremental pragmatics may lead to the language model being compelled to produce longer utterances – in section (4), where we apply the model to real-world data, in the form of the TUNA corpus.

However, this effect does not obtain in Spanish,

where adjectives are post-nominal. In the Spanish case, let our utterances be *vestido*, *vestido rojo*, *sombrero*, and *sombrero azul*, with the same referents and costs as before. Then there is no difference between the global and incremental models:

$$\begin{aligned} S_1^{\text{UTT-GP}}(\textit{vestido}|\text{R1}) &= 0.73 > \\ S_1^{\text{UTT-GP}}(\textit{vestido rojo}|\text{R1}) &= 0.27 \\ S_1^{\text{UTT-IP}}(\textit{vestido}|\text{R1}) &= 0.73 > \\ S_1^{\text{UTT-IP}}(\textit{vestido rojo}|\text{R1}) &= 0.27 \end{aligned}$$

When choosing the word to follow *vestido*, the incremental pragmatic speaker has no need to say *rojo* rather than STOP, since the goal of communicating the referent has already been completed by *vestido*. As a result, the speaker chooses the less costly option, STOP. The relevant difference here from the English case is that it is grammatical to stop after the first word (since the first word is a noun, not an adjective as in English).

A qualitative property which this example illustrates is a dislike in $S_1^{\text{UTT-IP}}$ for utterances which begin with a sequence of words which would mislead the incremental literal listener L_0^{WORD} . This is the basis on which anticipatory implicatures are formed, as discussed in section 3.2. This is not a hard constraint: utterances which would initially mislead an incremental listener are not categorically ruled out. However, the question of whether this behavior is empirically justified is a worthwhile topic for future investigation.

3.2 Anticipatory Implicatures

We now turn to a phenomenon concerning language interpretation, which we will model with an RSA listener, L_1^{WORD} .

Sedivy (2007) provides compelling empirical evidence that humans draw pragmatic inferences partway through utterances. For instance, when shown a scene with a tall cup, a tall pitcher, a short cup, and a key, a listener who hears “Give me the tall–” will fixate on the tall cup before the utterance is complete.

We take this as evidence for an incremental pragmatic listener L_1^{WORD} which can calculate an implicature by reasoning that, had the speaker intended to refer to the pitcher, they would not have had any motivation to say “tall”. By contrast, on the assumption that the speaker’s referent is the tall cup, the contrastive modifier serves to distinguish the intended referent from the short cup.

To model this implicature formally, we make the simplifying assumption that the possible utterances are *tall cup*, *short cup*, *tall pitcher*, *cup*, *pitcher*, and *key*. For consistency with the previous example, we assume the additive *cost* function from section 3.1.

On hearing *tall* as the first word of an utterance, L_1^{WORD} , the incremental pragmatic listener, can draw the following inference: the intended referent is likely to have been the tall cup, since had it been the tall pitcher, there would have been no need to use the contrastive modifier *tall*: $L_1^{\text{WORD}}(\text{the pitcher}|c = \square, \text{tall}) = 0.4$ while $L_1^{\text{WORD}}(\text{the tall glass}|c = \square, \text{tall}) = 0.6$.

This implicature is cancelable, and indeed, were the next word to be *pitcher*, we would exclude all referents but the pitcher. In this respect, the model represents the confusion created by uttering “tall pitcher”, where after the first word of the utterance, the majority of probability mass is on a referent (*the tall cup*) which, after the second word, has no probability mass.

4 Experimental Validation with TUNA

In order to observe the behavior of our incremental pragmatic model on real data, we make use of the TUNA corpus (van Deemter et al., 2006). TUNA is built around a referring expression task grounded in images. The images are coded using a fixed set of attributes, and the human-produced utterances are coded using the same attributes. Thus, TUNA lets us study the core content of naturally produced referring expressions without forcing us to confront the full complexity of natural language.

Our goal is to show that, when a cost is imposed which prefers shorter utterances, the incremental model $S_1^{\text{UTT-IP}}$ is less affected, and on average produces more two-word utterances than $S_1^{\text{UTT-GP}}$.

We hypothesize this on the basis of the preference of $S_1^{\text{UTT-IP}}$ for utterances where the choice of each word is made with high certainty. This means that informative one-word utterances which have reasonable probability of being extended with a second word will score lower than two-word utterances where the choice of the first word all but fixes the choice of the second. Since most one-word utterances admit the possibility of an extension to a second word, this dynamic would result in a preference for the longer, two-word utterances. This would provide further evidence that




Figure 4: An example target entity from the *furniture* domain, along with its coding as a dictionary, and the human generated referring expression for it in a context of other images.

the dynamic described in section 3.1 generalizes from an idealized example to real language.

4.1 Data

The TUNA corpus defines a reference game in the sense of figure 1a. Each *trial* contains a set of images (entities), of which one or several are the target, and a human-generated referring expression for the target in the context of all the images. We refer to the full set of target and non-target entities as the *context set*.

Both images and utterances are coded as sets of attributes (figure 4). This coding defines a semantics. For instance, in figure 4, the utterance “the grey desk” is true of the entity, since `type:desk` and `colour:grey` are included in its attributes. For the *furniture* domain, attributes such as color, object type, and size are coded. The *people* domain is more complex, coding for more attributes, including age, clothing, hair color, glasses, and orientation. Both domains also code for the position of the image relative to the other images in the context set.

4.2 Methods

For simplicity, we restrict our model to the *furniture* and *people* domains where only a single referent is provided, and consider only utterances of two words or fewer. These constitute 32% of the total utterances in the single referent corpora, and to our knowledge are not distinct in other ways than their length.

For each trial, the possible utterances are those from the set of all two-word utterances across the

entire corpus (either of furniture or people) which are compatible with at least one of the entities in the trial. We predict the set of optimal utterances (since there may be more than one utterance with maximum probability) for both $S_1^{\text{UTT-GP}}$ and $S_1^{\text{UTT-IP}}$. For our cost function, we assume all words have a cost of 1.0 except the STOP token, which has cost 0.0. Utterances cost the sum of their words. This has the effect of penalizing longer utterances.

For each trial, we have a set of entities as referents, with the designated target identified among these entities. In addition, we can define the set of all possible true utterances for a given trial. Thus, it is possible to make predictions according to both $S_1^{\text{UTT-GP}}$ and $S_1^{\text{UTT-IP}}$ for each trial without having to enrich the TUNA dataset in any way.

4.3 Results

As expected, we find a preference for longer utterances; out of the 114 *people* trials, $S_1^{\text{UTT-GP}}$ identifies 120 two-word utterances as optimal, compared to 287 for $S_1^{\text{UTT-IP}}$. In the 83 trials of the *furniture* domain, $S_1^{\text{UTT-GP}}$ marks 88 two word utterances as optimal, compared to 149 for $S_1^{\text{UTT-IP}}$. (More than one utterance may be optimal for a given trial, in the event that multiple utterances have the same, maximal probability of being chosen.)

An example of a representative case is the trial where the entity in figure 4 is the target, and no other distractors are grey, although others are desks. In this case, both “grey” and “a grey desk” are fully informative, in the sense of only being compatible with the target. With the cost term having the effect of penalizing longer utterances, $S_1^{\text{UTT-GP}}$ chooses “grey” as optimal. For $S_1^{\text{UTT-IP}}$, however, neither of these utterances are optimal, because probability is divided between stopping after “grey” and continuing with “desk”. Instead, the optimal utterance, “right middle”, describes the position of the target among the images of the context set. “Right” is not an available full utterance (as it is not attested in the data) and so no probability mass is lost by being divided between stopping and continuing with “middle”.

While this result offers a possible motivation for over-informative behavior, the nature of the relation is clearly nuanced – two-word utterances are not always more informative than one-word utterances – and merits further work. In particular, it

would be desirable to use a more direct proxy for over-informativity than preference for longer utterances.

5 Conclusion

In summary, we have defined a formal notion of incremental pragmatics, with respect to both production and comprehension, and shown that it differs in meaningful ways from global pragmatic behavior, at least within the RSA paradigm. The core differences are the sensitivity to word order, which varies cross-linguistically in a way which core pragmatic reasoning does not, and the ability of an incremental model to calculate implicatures partway through an utterance. We compared these models using simulations, and we assessed them using existing psycholinguistic data and using a new experiment with the TUNA corpus.

Exploration of RSA as a machine learning model is now underway, and this is helping to show the value of pragmatic reasoning in important NLP tasks. Such work forces us to confront the fact that the idealized RSA speaker agent must reason about all possible utterances – the normalization constant in (3) demands this. Prior work has sought to get around this by simplifying the space of possible utterances (Monroe and Potts, 2015) or by sampling a small set of utterances to approximate this normalization (Andreas and Klein, 2016; Monroe et al., 2017). Neither solution is ideal. The incremental approach offers a scalable alternative, as long as the incremental semantics is learned, as in Cohn-Gordon et al. 2018 and Vedantam et al. 2017.

There is also much to be done in assessing the incremental model in the context of on-line sentence processing. We have begun to identify the key properties of the model, but the degree to which these properties accord with empirical data on production and comprehension remains a largely open question.

6 Acknowledgments

Many thanks to the insightful comments of the reviewers, regarding both the presentation of the material and the substantive technical comments regarding the model itself.

This material is based in part upon work supported by the Stanford Data Science Initiative and by the NSF under Grant No. BCS-1456077.

References

- Jacob Andreas and Dan Klein. 2016. Reasoning about pragmatics with neural listeners and speakers. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1173–1182. Association for Computational Linguistics.
- Leon Bergen, Roger Levy, and Noah D. Goodman. 2016. Pragmatic reasoning through semantic inference. *Semantics and Pragmatics*, 9(20).
- Reuben Cohn-Gordon, Noah Goodman, and Christopher Potts. 2018. Pragmatically informative image captioning with character-level inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 439–443. Association for Computational Linguistics.
- Kees van Deemter, Ielka van der Sluis, and Albert Gatt. 2006. Building a semantically transparent corpus for the generation of referring expressions. In *Proceedings of the Fourth International Natural Language Generation Conference*, pages 130–132. Association for Computational Linguistics.
- Paul E Engelhardt, Karl GD Bailey, and Fernanda Ferreira. 2006. Do speakers and listeners observe the gricean maxim of quantity? *Journal of Memory and Language*, 54(4):554–573.
- Michael C. Frank, Andrés Gómez Emilsson, Benjamin Peloquin, Noah D. Goodman, and Christopher Potts. 2016. Rational speech act models of pragmatic reasoning in reference games. Ms., Stanford University.
- Michael C. Frank and Noah D. Goodman. 2012. Predicting pragmatic reasoning in language games. *Science*, 336(6084):998.
- Michael Franke. 2009. *Signal to Act: Game Theory in Pragmatics*. ILLC Dissertation Series. Institute for Logic, Language and Computation, University of Amsterdam.
- Michael Franke and Gerhard Jäger. 2014. Pragmatic back-and-forth reasoning. In Salvatore Pistoia Reda, editor, *Pragmatics, Semantics and the Case of Scalar Implicatures*, pages 170–200. Palgrave Macmillan UK, London.
- Albert Gatt, Anja Belz, and Eric Kow. 2009. The TUNA-REG Challenge 2009: Overview and evaluation results. In *Proceedings of the 12th European Workshop on Natural Language Generation (ENLG 2009)*, pages 174–182, Athens, Greece. Association for Computational Linguistics.
- Noah D Goodman and Michael C Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11):818–829.
- Noah D. Goodman and Andreas Stuhlmüller. 2013. Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science*, 5(1):173–184.
- H. Paul Grice. 1975. Logic and conversation. In Peter Cole and Jerry Morgan, editors, *Syntax and Semantics*, volume 3: Speech Acts, pages 43–58. Academic Press, New York.
- Theo Herrmann and Werner Deutsch. 1976. *Psychologie der objektbenennung*. Huber.
- Gerhard Jäger. 2007. Game dynamics connects semantics and pragmatics. In Ahti-Veikko Pietarinen, editor, *Game Theory and Linguistic Meaning*, pages 89–102. Elsevier, Amsterdam.
- Gerhard Jäger. 2012. Game theory in semantics and pragmatics. In Claudia Maienborn, Klaus von Heusinger, and Paul Portner, editors, *Semantics: An International Handbook of Natural Language Meaning*, volume 3, pages 2487–2425. Mouton de Gruyter, Berlin.
- Justine T. Kao, Jean Y. Wu, Leon Bergen, and Noah D. Goodman. 2014. Nonliteral understanding of number words. *Proceedings of the National Academy of Sciences*, 111(33):12002–12007.
- Will Monroe, Robert X. D. Hawkins, Noah D. Goodman, and Christopher Potts. 2017. Colors in context: A pragmatic neural model for grounded language understanding. *Transactions of the Association for Computational Linguistics*, 5:325–338.
- Will Monroe and Christopher Potts. 2015. Learning in the Rational Speech Acts model. In *Proceedings of 20th Amsterdam Colloquium*, Amsterdam. ILLC.
- Christopher Potts, 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.
- Paula Rubio-Fernández. 2016. How redundant are redundant color adjectives? an efficiency-based analysis of color overspecification. *Frontiers in psychology*, 7:153.
- Julie C Sedivy. 2007. Implicature during real time conversation: A view from language processing research. *Philosophy compass*, 2(3):475–496.
- Ramakrishna Vedantam, Samy Bengio, Kevin Murphy, Devi Parikh, and Gal Chechik. 2017. Context-aware captions from context-agnostic supervision. In *Computer Vision and Pattern Recognition (CVPR)*, volume 3.
- Adam Vogel, Max Bodoia, Christopher Potts, and Dan Jurafsky. 2013. Emergence of Gricean maxims from multi-agent decision theory. In *Human Language Technologies: The 2013 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 1072–1081, Stroudsburg, PA. Association for Computational Linguistics.

Erica J. Yoon, Michael Henry Tessler, Noah D. Goodman, and Michael C. Frank. 2016. Talking with tact: Polite language as a balance between kindness and informativity. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, pages 2771–2776. Cognitive Science Society.