

Pragmatic language interpretation as probabilistic inference

Noah D. Goodman^{a,*}, Michael C. Frank^a

^a*Department of Psychology, Stanford University, 450 Serra Mall, Stanford, CA 94305*

Abstract

Understanding language is more than the use of fixed conventions and more than decoding combinatorial structure. Instead, comprehenders make exquisitely sensitive inferences about what utterances mean given their knowledge of the speaker, the language, and the context. Building on developments in game theory and probabilistic modeling, we describe the rational speech act (RSA) framework for pragmatic reasoning. RSA models provide a principled way to formalize inferences about meaning in context; they have been used to make successful quantitative predictions about human behavior in a wide variety of different tasks and situations, and they explain why complex phenomena like hyperbole and vagueness occur. More generally, they provide a computational framework for integrating linguistic structure, world knowledge, and context in pragmatic language understanding.

... one of my avowed aims is to see talking as a special case or variety of purposive, indeed rational, behavior ...

Grice (1975), p. 47

1. Understanding language

Language is central to the successes of our species; with language we can coordinate our actions, learn from each other, and convey our innermost thoughts. From sounds to syntax, natural languages provide structured methods of combining discrete materials to generate an infinite variety of sentences. Yet this

*Corresponding author, please send correspondence to ngoodman@stanford.edu.

discrete combinatorics does not fully explain how speakers can use language so flexibly to achieve social goals. The interpretation of a particular utterance can itself be almost infinitely variable, depending on factors such as the identity of the speaker, the physical context of its use, and the previous discourse. While the systematization of structural features of language is one of the proudest accomplishments of cognitive science [e.g., 15, 40, 30], its contextual flexibility – its pragmatics – has been stubbornly difficult to formalize.

Grice [34] presented an initial framework theory for pragmatic reasoning, positing that speakers are taken to be cooperative, choosing their utterances to convey particular meanings. Gricean listeners then attempt to infer the speaker’s intended communicative goal, working backwards from the form of the utterance. This goal inference framework for communication has been immensely influential [e.g., 37, 71, 17, 53]. But attempts to build on these ideas by providing a specific set of formal principles that allow the derivation of pragmatic inferences have met with difficulty.

For example, the core of Grice’s proposal was a set of **conversational maxims** (see Glossary). Inferences about speakers’ behavior relative to these maxims (be truthful, relevant, informative, and perspicuous) could lead to **implicatures** – inferences about their intended meaning. Formalization of the Gricean notion of implicature using the maxims is difficult, however, and many post-Gricean theories have instead proposed alternative sets of principles [71, 53]. An important test of the difficulty of this theoretical project is that the burgeoning experimental psycholinguistic literature attempting to measure pragmatic inference has found these principles only modestly useful [11, 39, 60]. In addition, this sort of informal theory of pragmatics can make only directional, qualitative predictions with respect to experimental data that are typically graded and quantitative.

An alternative strand of Gricean thought has had more success in making contact with data. Grice’s core insight was that language use is a form of rational action; thus, technical tools for reasoning about rational action should elucidate linguistic phenomena. Such a goal-directed view of language production has led to the development of engineering systems for natural language generation [18] that have in turn been applied as theories of human language production [e.g., 75]. Concurrently, the tools of game-theory – which allow for the characterization of rational actions with respect to defined utilities – have provided a vocabulary for formal descriptions of pragmatic phenomena [e.g., 6, 41]. The recent work we focus on here builds on these developments, combining them with a more detailed view of cognition that arises from the Bayesian cognitive modeling tradition.

Probabilistic, or Bayesian, models have been at the core of a set of recent attempts to understanding the interplay between structured representations and graded or statistical information [74]. These models have been an important tool for understanding non-linguistic varieties of rational action, integrating belief understanding with action planning [2]. A critical feature of these models

is that they use the probability calculus to describe inferences under uncertainty. Within formal models of pragmatics, this uncertainty stems from a variety of sources, including uncertainty about speakers’ goals and beliefs, uncertainty about the discourse and broader context, and even uncertainty about the meanings of words.

In the remainder of this paper we describe the probabilistic approach to pragmatics. We begin by presenting the **rational speech act (RSA) model**, and the growing body of empirical data supporting its utility in explaining pragmatic reasoning. We next discuss extensions to RSA that allow it to be applied to non-literal uses of language like hyperbole, irony and metaphor, to cases of vagueness and ambiguity, and to complex interactions between pragmatics and compositional syntax/semantics. We close by considering the broader applications of – and challenges for – probabilistic pragmatics models.

2. A “rational speech act” model

The “rational speech act” (RSA) model implements a social cognition approach to utterance understanding. At its core, it captures the idea (due to Grice, David Lewis, and others) that speakers are assumed to produce utterances to be helpful yet parsimonious, relative to some particular topic or goal. Listeners then understand utterances by inferring what such a helpful speaker must have meant, given what she said.¹ The first of these basic assumptions is formalized by viewing the speaker as a utility-maximizing agent (where the effort of language production is costly, but communicating information is beneficial). The listener then updates his beliefs via Bayesian inference.

The pragmatic listener infers the state of the world, w , using Bayes’ rule, given the observation that the speaker chose a particular utterance, u :

$$P_L(w | u) \propto P_S(u | w)P(w)$$

The key assumption he must make is that the speaker is *approximately rational*, that is, that she has chosen her utterances in proportion to the utility she expects to gain.

$$P_S(u | w) \propto \exp(\alpha U(u; w))$$

The speaker chooses u from a set of *alternative utterances* (see Outstanding Questions). The parameter α captures the extent to which the speaker maximizes her utility – how rational she will be. The basic speaker utility used in RSA captures the social benefit of providing epistemic help to a listener:

¹For clarity throughout, we use a female pronoun for Alice, the speaker, and a male pronoun for Bob, the listener.

$$U(u; w) = \log P_{\text{Lit}}(w \mid u)$$

This expression measures how certain the listener becomes about the intended world, after hearing the utterance; in order to avoid an infinite recursion and provide an entry point for conventional (semantic) meaning, the speaker is assumed to consider a simpler listener, the *literal listener* P_{Lit} . The literal listener again updates his beliefs in accord with Bayesian inference, under the assumption that the literal meaning of the utterance is true:

$$P_{\text{Lit}}(w \mid u) \propto \delta_{\llbracket u \rrbracket(w)} P(w)$$

This definition of the literal listener requires a semantic denotation for each sentence, $\llbracket u \rrbracket$, in which a sentence has the value true or false when applied to a particular state of affairs, w (the “world”). This denotation is how conventional meaning enters the pragmatic reasoning process, and, as we describe briefly below, it connects RSA to work in lexical and compositional semantics [36, 22].

Consider the scenario in Figure 1, in which speaker and listener share a world of three faces – one with hat and glasses (HG), one with glasses only (G), and one with neither (N); one of these (known to the speaker but not the listener) is the “friend.” The speaker says “my friend has glasses,” presupposing that there is a single friend. Experimental participants, who know that the only alternative utterance was “my friend has a hat,” tend to share the intuition that this sentence refers to G and not HG or N [72, 73]. While real-world language has many more utterances available (e.g., “my friend has glasses, but no hat”), this restricted scenario serves to illustrate the underlying dynamics of pragmatic reasoning.

Under RSA, listener L reasons about S (a simplified internal representation of the speaker), who in turn reasons about Lit (a yet more simplified internal model of the listener). Lit updates his beliefs based on a straightforward denotation: “glasses” applies to both G and HG, while “hat” applies only to HG. Thus $P_{\text{Lit}}(w \mid \text{“hat”})$ places all probability on the friend being HG, while $P_{\text{Lit}}(w \mid \text{“glasses”})$ places equal probability on G and HG (see innermost thought bubbles in Figure 1). The speaker S who intends to communicate that HG is the friend will thus tend to choose the more informative “hat”; but if she intends to communicate that G is the friend, she will use “glasses.” Finally, upon hearing “glasses,” the listener L infers that this likely refers to G (reflecting the counterfactual that if S had been talking about HG, she would have said “hat” instead).

In the simplest RSA model, as illustrated above, the speaker values providing epistemic help – information – to the listener. But the model can also be extended to create a more sophisticated speaker who is uncertain about the world state, who avoids costly utterances, or who aims to provide relevant information.

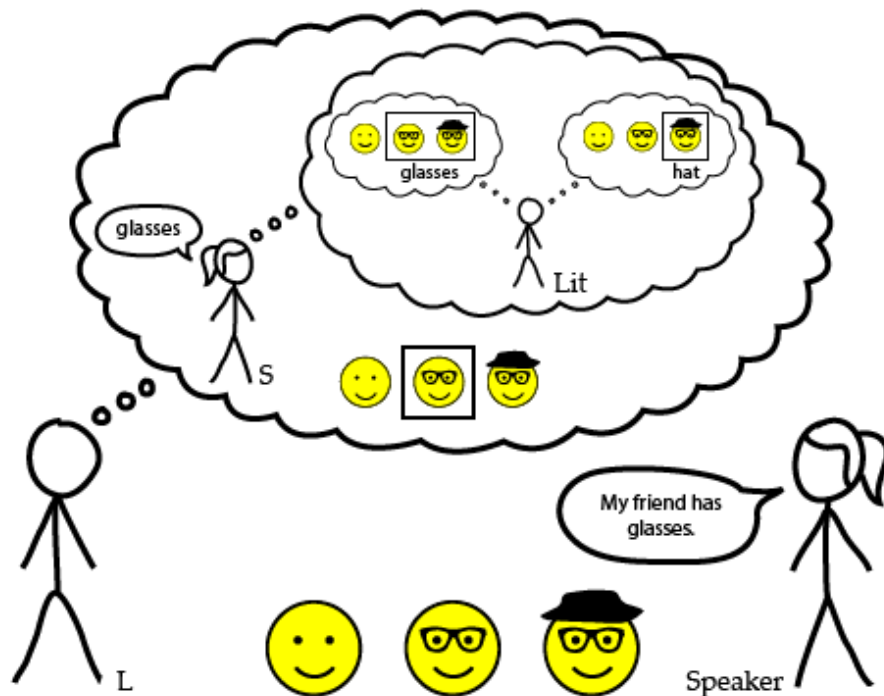


Figure 1: Application of RSA-style reasoning to a signaling game (shown by the three faces along the bottom). Agents are depicted as reasoning recursively about one another’s beliefs: listener L reasons about an internal representation of a speaker S, who in turn is modeled as reasoning about a simplified literal listener, Lit. Boxes around targets in the reference game denote interpretations available to a particular agent.

Connections to other theoretical approaches and aspects of language then become straightforward. For instance, by modifying the speaker’s utility function, we can model the notion of *topic-relevant* information, which connects to linguistic ideas about the “question under discussion” [68]. As a second example, RSA can be combined with the noisy channel approach to language comprehension [54], in order to explain the communicative use of sentence fragments and prosodic stress [7].

In sum, RSA models replace Grice’s maxims with a single, utility-theoretic version of the cooperative principle [27]. This formulation is based on utilities that can reflect the communicative and social priorities of a complex, real-world agent.

3. Empirical support for RSA

The example shown above in Figure 1 is an instance of a signaling game of the type initially introduced by Lewis [55]. Such games are a valuable tool for exploring pragmatic inferences in context, and experiments testing the RSA framework have used games of this type to make quantitative measurements of a variety of different inferences. For example, one study [24] used a one-shot, web-based paradigm to present participants with geometric shapes in a variety of different configurations. Using a betting paradigm (participants were asked to distribute \$100 between response options), these experiments collected separate judgments about what a speaker would say, a listener would interpret, and about baseline expectations for reference (corresponding to the prior $P(w)$). The RSA model showed a tight fit to listeners’ aggregate judgments when combined with empirical measurements of the prior distribution: P_S and P_L models correlated very strongly with participants’ average bets on what to say and how to interpret, respectively.

Although in this initial work RSA was used to simulate the behavior of both speakers and listeners, most subsequent work has focused on the behavior of listeners alone. This work follows the idea that RSA captures listeners’ (perhaps optimistic) assumptions about the rational behavior of speakers. Thus, RSA is ‘rational’ in the sense of assuming that speakers are rational; a separate question is how rational speakers actually are (see **Box 11**). In addition, though most research using RSA models has focused on mature language comprehenders, these models have also been the inspiration for a variety of developmental work (see **Box 10**).

A variety of other work has replicated and extended the initial findings using similar signaling-game paradigms. A tight replication of the initial results [66] reproduced the basic findings and explored a set of variants to the initial RSA utility function. Other studies [13] found that RSA predicted judgments in a communication game using much more complex spatial language stimuli, albeit with somewhat noisier fits. Thus, RSA with an epistemic utility can predict judgments in simple signaling games across variations in both participant sample and stimulus.

One question raised by this initial work was the level of **social recursion** (see Glossary) that best fit human performance. The presentation of RSA given above is stated in terms of a minimal recursion (a listener reasons about speaker, who in turn reasons about a literal listener) but much greater depths of reasoning are in principle possible. The evidence is mixed on whether deeper levels of recursion are commonly seen in language comprehension. In a variety of experiments exploring this issue, participants tended to show chance-level performance for signaling systems that required deeper levels of recursion to find unique interpretations [72, 19, 77]. More recent studies [26], however, showed some evidence of deeper recursion for a subpopulation of participants (approx. 15%) in a more complex paradigm, consistent with work on competitive economic games where

deeper recursions are sometimes found [12]. This heterogeneity – and its dependence on individual and contextual differences – is an interesting topic for future work.

A number of other studies have tested RSA with more elaborated utility functions (see **Box on Refinements to the speakers utility**). For instance, a speaker might be expected to produce a less informative utterance when the more informative one is much harder to say. This tendency can be formalized by including a cost term in the speaker’s utility; with this modification, RSA predicts the impact of production costs on interpretations of the listener. Work exploring this extension [8] showed that participants in a reference game are indeed sensitive to the cost of message choices: the effect of alternative possible messages on a listener’s inferences is modulated by their cost, in dollars. Related studies [20] tested the effect of production difficulty by manipulating how quickly the speaker could type on an on-screen keyboard; participants’ interpretations reflected this difficulty as predicted. Additional work has used proxies for production cost such as number of words and their frequencies in explorations of phenomena such as negation [58] and the choice of noun used to refer to an object [33].

Finally, in addition to ad-hoc signaling systems, RSA provides a way to describe reasoning about classic linguistic implicatures. Perhaps the best-studied of these is the scalar implicature that “some of the letters had checks inside” implicates that not all did. One study [32] measured participants’ judgments about the interpretations of quantifiers and number words in exactly this situation and found that these judgments were well-predicted by RSA. In addition, a critical feature of this study was the inclusion of an epistemic manipulation (e.g., whether all of the letters had already been opened). By using expected informativity (see **Box**) to account for the speaker’s limited perceptual access, the model was able to predict differing patterns of listener judgments based on different levels of speaker uncertainty. These empirical findings are congruent with other recent demonstrations of the importance of epistemic reasoning in pragmatic implicature [8, 10], and the theoretical account accords with other probabilistic treatments of scalar implicature [69]. They also highlight the way the RSA framework provides a (non-modular) theory for interactions between language and non-linguistic cognition. We next turn to a variety of other extensions to the basic RSA model that explore additional interactions.

4. Uncertainty about the speaker: Joint reasoning

In the basic RSA model, the listener has a specific model in mind of how a speaker will behave. But what should a listener do if he is not sure what speaker model is appropriate? Speakers can differ in knowledge, communicative goals, and many other aspects; these differences can lead a listener to arrive at different interpretations of the same utterance. Recent work has addressed this issue by positing a joint inference: What type of speaker am I interacting with

and what is the world like, given the utterance I heard? Formally this *uncertain* RSA (or uRSA) framework requires only a small change:

$$P_L(w, s | u) \propto P_S(u | w, s)P(s)P(w)$$

where the new variable s parametrizes different speaker types. In practice, s can refer to any factor that might influence the speaker’s behavior, including uncertainty about conversational topic, word meanings, background knowledge, or general discourse context. This modification allows uRSA to capture a much wider variety of linguistic phenomena; intuitively, an uRSA listener is a more realistic cognitive agent than the RSA listener, who was restricted to the specifics of a very particular context and goal. To illustrate this intuition, we provide three examples of phenomena captured by uRSA (but not by basic RSA): non-literal language, vagueness, and embedded implicatures.

Non-literal or figurative language – utterances that are easily interpreted but not “actually true” – poses a problem for nearly all formal models of language understanding. How can tropes like hyperbole, sarcasm, and metaphor be interpreted, and why are they used? Under uRSA, these uses can be described as arising from uncertainty about the topic of conversation. If the speaker is expected to provide information relevant for a particular topic, the pragmatic listener will only update his beliefs along this topical dimension. Within uRSA, the interaction between uncertainty about the speaker’s intended topic and her intended meaning about that topic can drive complex interpretations.

Hyperbolic utterances such as “the electric kettle cost \$1,000” are a key case study [45]. In this example, the number \$1,000 can be interpreted as conveying information about the speaker’s affect, not the actual price, in part because one thousand dollars is an implausibly high price for a kettle. As shown in Figure 2, the uRSA model captures this intuition by positing that the topic of the speaker’s utterance may be the actual price of the kettle, the speaker’s opinion about the price, or some combination of the two. Since the listener does not know the topic, he jointly infers it together with the likely true price of the kettle and the speaker’s affect. When the uttered price is implausible it becomes more likely that the speaker is aiming to convey her opinion, and using \$1,000 (a price which most people would find too high) to do so. In this way, the listener’s joint inference can yield a non-standard topic and hence a non-literal interpretation.

By extending the space of affect to include both valence and arousal, the same model predicts verbal irony [44]. And a similar approach has been suggested for simple metaphors [43], such as “John is a shark.” Here the potential topics include not affect, but features of the target, such as how vicious John is and how likely he is to swim underwater. In each of these cases of figurative language, the uRSA model accounts for almost all of the explainable variance in human interpretations, a striking result considering the complexity and subtlety of these phenomena.




Figure 2: Uncertain RSA-style reasoning applied to hyperbole. Listener L reasons jointly about the price of the item and the speaker’s affect. In doing so he considers two speakers, one who is primarily interested in conveying her affective response to the kettle, and one who is primarily interested in conveying the actual price. (The full model also considers speakers, not pictured, that wish to convey approximate price and combinations of these goals.) Each of these speakers is modeled as reasoning about a literal listener who interprets the utterance literally (indicated by the box selecting the “\$1000” state), but focus on different aspects of the situation (price on the left and affect on the right).

Many linguistic descriptions, especially adjectives, are both context-sensitive and vague. Providing precise definitions for words like *expensive* or *tall* has been a persistent challenge for philosophers and semanticists [79]. uRSA models address this challenge by assuming that word meanings can differ between speakers and contexts, and that these meanings themselves can be a subject for inference. In the case of scalar adjectives like *tall*, the uncertainty is over the threshold required: what height is required before an object counts as tall? Under the uRSA model, judgments about meaning take into account two conflicting pressures: On one hand, a stricter threshold for tallness makes the term *tall* more informative. For instance “Bob is tall” tells us a great deal if *tall* requires a height greater than eight feet. But, on the other hand, a stricter threshold makes such a sentence quite unlikely to be true *a priori*. By negotiating this balance between informativity and plausibility, uRSA accounts for three key phenomena of vague adjectives [51]: the inferred meaning depends on

the class (tall for a tree vs. tall for a person), there are borderline cases, and the interpretations are subject to a *sorites* paradox (no single minimal increment in height will make you tall, but enough increments will). The process of reasoning about meaning that are modeled by uRSA might even interact with learning processes to produce more long-lasting inferences about word meaning, leading to language change (see **Box on Language use and language change**).

Finally, this uRSA approach allows for progress on an important puzzle in recent discussions of pragmatic inference: embedded implicature [29, 14]. Embedded implicatures occur when quantifiers are nested within one another, as in sentences like “Exactly one letter is connected with some of its circles.” In these cases, some experimental evidence suggests that participants access the interpretation that one letter is connected with some *but not all* of its circles, an interpretation that standard Gricean theories cannot generate [14]. Recent work [65] has replicated these interpretations in a series of large-scale experiments and confirmed that basic RSA models could not capture them. An implementation of uRSA that jointly infers word meanings and world state [9, 65] showed a good fit to the overall pattern of data. Sentence meanings in this model are built by composing uncertain word meanings, showing how uRSA is a fruitful way to incorporate pragmatic reasoning into compositional semantic systems.

5. Concluding remarks and future directions

Context-dependence is one of the core features of natural language. Yet because of the informal nature of theorizing about this context-dependence, pragmatics has often been treated as a theoretical “wastebasket,” in which unexplained phenomena are hidden [3]. Countering this trend, new formal theories of pragmatics make quantitative predictions about a wide variety of phenomena that have previously been considered too difficult to operationalize. These include implicature, vagueness, non-literal language, and the myriad other cases where linguistic meaning is changed by context.

The key tool in this work is the Rational Speech Act framework, which builds upon and synthesizes a number of formal traditions in the study of human inference, from game theory to models of human reasoning. The RSA approach also builds on existing work on semantic representation – using a compositional semantics *a la* Montague [22] – and contributes back to semantics, providing a specific mechanism by which underspecified meanings become precise, in context. Rather than formalizing only a single hypothesis about pragmatic language understanding, RSA provides a framework in which many variations can be explored. Varying assumptions about the speaker’s utility (see **Box: Refinements to the speaker’s utility**) and listener’s uncertainty (see Section 4), for instance, yields a spectrum of hypotheses that can be evaluated against quantitative experimental data.

While it has been successful in many recent cases, it may emerge that the RSA approach is not able to capture some aspects of language understanding,

either because the foundational, Bayesian, tools it relies on are inadequate [42], or because pragmatic effects arise from sources not easily incorporated into RSA. Optimistically, however, RSA can be combined with other approaches when needed. For instance, the alternative utterances in RSA can be restricted [65] using previously-proposed grammatical mechanisms [23]. And increasingly, methods in machine learning have been used to supplement RSA with powerful learning mechanisms [57]. Indeed, this cross-fertilization is among the most encouraging outcomes of work on RSA.

The RSA framework is a *computational* level description of the language user’s competence, in Marr’s sense [56]. There are many possible ways a cognitive agent could implement RSA at the algorithmic level, and it is unclear which might match the speed and competence of human language understanding and production. These alternatives must further be evaluated for their ability to explain the processing signatures of language comprehension such as reaction times and eye gaze [21, 58]. Yet even as a computational-level framework, RSA inspires different intuitions about processing than previous theories. For example, RSA-style reasoning makes pragmatic inferences a fundamental part of language comprehension, in which the ultimate goal of all interpretation is to settle on the intended meaning, given both the literal semantics of the utterance and the broader pragmatic context. This framing contrasts with Gricean analyses, in which pragmatics enters when the violation of a maxim leads to reasoning to “repair” the interpretation, and correspondingly slower processing, a view that has been challenged both theoretically and empirically [e.g., 53, 35].

Future extensions of RSA will likely include worlds with more and richer structure; a thorough and practical theory of pragmatic alternatives; more sophisticated discourses that unfold over many utterances; and utility structures that better take into account the complexities of social interaction. On the practical side, computing the predictions of RSA models can become prohibitive when the number of world states or utterances grows large. Further development of algorithms to implement RSA is needed. These developments may go together with new algorithms for learning aspects of the underlying semantics, which will open up new applications for the RSA approach in computational linguistics and artificial intelligence [31, 76, 57, 1, 61].

The work outlined in this review represents steps towards a comprehensive, formal theory of language understanding in context. Although much further work will be required, RSA models and their uRSA extensions have proven to be useful tools for explaining both qualitative and quantitative empirical data across a wide range of tasks and contexts. Language is central to the human experience. We hope our work sheds light on how its structure and systematicity can still give rise to such an astonishingly flexible communication system.

6. Acknowledgements

We gratefully acknowledge funding from a James S. McDonnell Foundation Scholar Award to Goodman, ONR Grants N00014-13-1-0788 and N00014-13-1-0287 to Goodman and Frank, and NSF BCS Grant 1456077 to Frank.

7. Trends

- Rational speech act (RSA) models provide a quantitative framework to capture intuitions about pragmatic reasoning in language understanding.
- Extensions to RSA that allow for reasoning about the speaker – for instance, her goals and word usage – can capture many otherwise puzzling phenomena including vagueness, embedded implicatures, hyperbole, irony, and metaphor.
- The RSA framework can inform psycholinguistic processing experiments, linguistic theory, and scalable natural language processing models.

8. Box: Outstanding questions

- *Social recursion.* How deeply do human comprehenders reason about others' intentions? Is depth of recursion – “I think that you think that...” – constant, or does it vary across situations?
- *Alternatives.* How are alternative utterances computed? Do they depend on the language grammar? On situational factors?
- *Linguistic goals.* How do “Gricean” utilities – the drive to be informative yet succinct – relate to other social goals like conveying affect, or establishing relationships? How do cooperative and competitive goals mix in language use?
- *Dialogue.* How can RSA be used to model the evolution over the course of a conversation of a partners' utilities, possible goals, and the context more broadly?
- *Learning and language change.* How do pragmatic language understanding and language learning interact? How and when does pragmatic language use lead to language change?
- *Algorithmic challenges.* Given the potential complexity of recursive pragmatic computations, how is language processed so quickly? How can RSA models be “scaled up” for natural language processing tasks?

9. Box: Refinements to the speaker’s utility

The notion of the speaker’s utility – what is rewarding for a speaker – is central to the RSA approach. The basic RSA model captures the speaker’s need to be *informative* to a listener: $U(u; w) = \log P_{\text{Lit}}(w | u)$. Different utilities lead to different kinds of speakers, which in turn lead to different interpretations by the pragmatic listener. Several utility refinements (and their combinations) have been considered in recent work:

- Utterance cost: In order to capture a tendency of speakers to be parsimonious we can simply add a cost term: $U(u; w) = \log P_{\text{Lit}}(w | u) + \text{cost}(u)$. The cost may reflect actual production cost (such as number of words) or proxies such as word frequency. This extension yields effects similar to Grice’s maxim of manner [9].
- Speaker uncertainty: When the speaker does not have full knowledge of the world she should choose an utterance according to *expected utility*: $U(u; k) = \mathbb{E}_{P(w|k)}[U(u; w)]$, where k summarizes the speaker’s knowledge or observations. This extension correctly predicts interactions between a speaker’s knowledge and a listener’s interpretations [32].
- Topic relevance: Although it may be highly informative to provide detailed descriptions, such detail is not always relevant. Relevance can be captured by introducing a topic of conversation [sometimes known as a *Question Under Discussion*, 68] and adjusting the epistemic utility to reflect only information about this topic: $U(u; w, t) = \log \sum_{w' \text{ s.t. } t(w')=t(w)} P_{\text{Lit}}(w' | u)$. Here the topic is encoded in a function t that takes a complete world and yields some subset or summary; for instance, in the case of hyperbole [45], t can pick out only the speaker’s affect, dropping objective states.
- Other social goals: Language is often used not just to inform, but also to flirt, insult, comfort, and to pursue myriad other social goals. For example, non-informational utilities, e.g. utility directed towards kindness, can produce behaviors that appear polite [81].

10. Box: RSA and children’s developing pragmatic competence

From a very early age, children are oriented towards communication, understanding the function of language for information transfer and repurposing their limited linguistic means to achieve a wide variety of ends [78, 16]. In light of this general early orientation, the literature on pragmatic development specifically has been puzzling: older children very reliably fail to make scalar implicatures under a range of circumstances [59]. In one striking example, five-year-olds endorsed the statement that “some of the horses jumped over the fence” even when three out of three of a set of horses had made the jump [62]. RSA-style

models can provide a framework for thinking about this disconnect between early communicative successes and later pragmatic failures.

Recently, theorists have proposed that children’s apparent difficulties with pragmatic implicatures may have resulted from their inadequate knowledge of relationships between lexical alternatives rather than difficulty with pragmatic computations more generally [4]. Congruent with this idea, three-year-old children show signs of successful implicature computations in the kinds of signaling games shown in Figure 1, where the referential alternatives are all simple objects that are visible in the scene [73]. And children in the same age range were able to use an implicature to guess the meaning of a novel word [25] or a novel context [38], showing the kind of inferential flexibility posited in RSA accounts.

These findings support the idea that even young children are able to make flexible pragmatic inferences, and are consistent with the application of RSA-style reasoning, albeit with limits on the available semantic alternatives. Future research will be required, however, to test whether RSA (or some capacity-limited modification) could make quantitative predictions about pragmatic development.

11. Box: Producing referring expressions

RSA stands for the “rational speech act” model, indicating that listeners idealize speakers as rational. Are speakers in fact rational in a meaningful way? And if so, how can this conclusion be integrated with the large body of evidence indicating that speakers are egocentric, error prone, and subject to idiosyncratic production preferences [47, 50, 28]?

Although our initial studies collected judgments about language production in extremely restricted tasks [24], most recent work using the RSA model has focused on modeling listeners’ judgments, rather than speakers’ productions. One reason for this choice is that often the most interesting pragmatic inferences come about when speakers are not maximally informative. For example, in the signaling game shown in Figure 1, helpful speakers will often overspecify and say “glasses and no hat” [5]. But this unnecessarily redundant utterance may in fact be a reasonable response to uncertainty about whether a conversational partner will in fact draw the desired implicature. More generally, speakers’ production choices are a promising area for future research using RSA models with a broader range of utility functions (see Box) and various sources of potential miscommunication (a topic of ongoing research).

Nevertheless, it is clear that in their natural behavior, speakers make production decisions under time pressure and a variety of cognitive demands [52]. Integrating these demands with the predictions of utility-theoretic models should be an important challenge for future work.

12. Box: Language use and language change

The pragmatic processes described by RSA models occur in the moment of communication, but can have a set of effects that ripple out through language as a whole. The construal of an individual communication event can influence learning processes, which in turn can lead to systematic changes in word meanings [70]. Words that are too narrow in their denotation can be pragmatically extended [45], while words that are too broad can be narrowed via implicature. Over time, word meanings may converge to the appropriate level of ambiguity to enable efficient communication [64]. In this sense, in-the-moment pragmatic interpretation may bootstrap long term language change.

The processes of change that promote efficient communication have been explored extensively within the iterated learning paradigm [48]. This framework can also be used to express the competing pressures of learnability and communication. When languages are selected only to be learnable, they often become degenerate, including only a single word [63]. But when they include a countervailing pragmatic pressure, which can be modeled via RSA, expressive and compositional languages can emerge [49].

If pressures for efficient communication lead to language change, then these pressures should be visible in the lexicons of human languages. Indeed, a recent body of evidence suggests that the typological distribution of languages in particular semantic domains reflects the range of optimal communication systems [e.g., 67, 46, 80]. An important future direction is to understand whether this typological distribution is predicted to arise from iterated learning with a population of RSA-like language users.

13. Glossary

- **Rational speech act model (RSA).** A class of probabilistic model that assumes that language comprehension in context arises via a process of recursive reasoning about what speakers would have said, given a set of communicative goals.
- **Uncertain RSA models (uRSA).** A specific extension of RSA models that allow for joint inferences about both the speaker’s intended meaning and other aspects of the interaction, such as the topic, the context, or even the meanings of particular words.
- **Implicature.** An inference about the meaning of an utterance in context that goes beyond its literal semantics. Implicatures are typically *cancellable* in that they can be contradicted, as in “Some of the students passed the exam; indeed, all of them did.”

- **Scalar implicature.** An implicature that arises when a speaker didn't use a stronger alternative term, leading to narrowing of the interpretation. For instance, hearing "some" usually leads to a scalar implicature that "not all."
- **Conversational maxims.** A set of principles described by Grice [34] as a theory of how listeners reason about speakers' intended meaning to arrive at pragmatic implicatures.
- **Social recursion.** Reasoning that involves two people, such as listener and speaker, thinking about each other: "I think that you think that I think that...." This may proceed to a finite depth, or continue *ad infinitum*.

References

- [1] Andreas, J., & Klein, D. (2016). Reasoning about pragmatics with neural listeners and speakers. In *EMNLP 2016*.
- [2] Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113*, 329–349.
- [3] Bar-Hillel, Y. (1971). Out of the pragmatic wastebasket. *Linguistic inquiry*, *2*, 401–407.
- [4] Barner, D., Brooks, N., & Bale, A. (2011). Accessing the unsaid: The role of scalar alternatives in children's pragmatic inference. *Cognition*, *118*, 84 – 93. doi:10.1016/j.cognition.2010.10.010.
- [5] Baumann, P., Clark, B., & Kaufmann, S. (2014). Overspecification and the cost of pragmatic reasoning about referring expressions. In *Proceedings of the Thirty-Sixth Annual Conference of the Cognitive Science Society*.
- [6] Benz, A., Jäger, G., & Van Rooij, R. (2006). An introduction to game theory for linguists. In *Game theory and pragmatics* (pp. 1–82). Springer.
- [7] Bergen, L., & Goodman, N. D. (2015). The strategic use of noise in pragmatic reasoning. *Topics in Cognitive Science*, .
- [8] Bergen, L., Goodman, N. D., & Levy, R. (2012). That's what she (could have) said: How alternative utterances affect language use. In *Proceedings of the 34th Annual Conference of the Cognitive Science Society*.
- [9] Bergen, L., Levy, R., & Goodman, N. D. (2016). Pragmatic reasoning through semantic inference. *Semantics & Pragmatics*, .
- [10] Breheny, R., Ferguson, H. J., & Katsos, N. (2013). Taking the epistemic step: Toward a model of on-line access to conversational implicatures. *Cognition*, *126*, 423–440.

- [11] Breheny, R., Katsos, N., & Williams, J. (2006). Are generalised scalar implicatures generated by default? an on-line investigation into the role of context in generating pragmatic inferences. *Cognition*, *100*, 434–463.
- [12] Camerer, C. F., Ho, T.-H., & Chong, J.-K. (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, (pp. 861–898).
- [13] Carstensen, A., Kon, E., & Regier, T. (2014). Testing a rational account of pragmatic reasoning: The case of spatial language. In *Proceedings of the 36th Annual Meeting of the Cognitive Science Society. Austin, TX: Cognitive Science Society* (pp. 2009–2013).
- [14] Chemla, E., & Spector, B. (2011). Experimental evidence for embedded scalar implicatures. *Journal of semantics*, (p. ffq023).
- [15] Chomsky, N. (1965). *Syntactic structures*. Walter de Gruyter.
- [16] Clark, E. V., & Amaral, P. M. (2010). Children build on pragmatic information in language acquisition. *Language and Linguistics Compass*, *4*, 445–457.
- [17] Clark, H. (1996). *Using Language*. Cambridge, UK: Cambridge University Press.
- [18] Dale, R., & Reiter, E. (1995). Computational interpretations of the gricean maxims in the generation of referring expressions. *Cognitive science*, *19*, 233–263.
- [19] Degen, J., & Franke, M. (2012). Optimal reasoning about referential expressions. In *Proceedings of SemDIAL*.
- [20] Degen, J., Franke, M., & Jäger, G. (2013). Cost-based pragmatic inference about referential expressions. In *Proceedings of the 35th annual conference of the cognitive science society* (pp. 376–381).
- [21] Degen, J., & Tanenhaus, M. K. (2015). Processing scalar implicature: A constraint-based approach. *Cognitive science*, *39*, 667–710.
- [22] Dowty, D. R., Wall, R., & Peters, S. (2012). *Introduction to Montague semantics* volume 11. Springer Science & Business Media.
- [23] Fox, D., & Katzir, R. (2011). On the characterization of alternatives. *Natural Language Semantics*, .
- [24] Frank, M., & Goodman, N. (2012). Predicting pragmatic reasoning in language games. *Science*, *336*, 998.
- [25] Frank, M. C., & Goodman, N. D. (2014). Inferring word meanings by assuming that speakers are informative. *Cognitive Psychology*, *75*, 80–96. URL: <http://langcog.stanford.edu/papers/FG-cogpsych2014.pdf>.

- [26] Franke, M., & Degen, J. (2016). Reasoning in reference games: Individual- vs. population-level probabilistic modeling. *PloS one*, *11*, e0154854.
- [27] Franke, M., & Jäger, G. (2016). Probabilistic pragmatics, or why bayes rule is probably important for pragmatics. *Zeitschrift für Sprachwissenschaft*, *35*, 3–44.
- [28] Gatt, A., van Gompel, R. P., van Deemter, K., & Krahmer, E. (2013). Are we bayesian referring expression generators. In *Proceedings of CogSci*. volume 35.
- [29] Geurts, B., & Pouscoulous, N. (2009). Embedded implicatures?!? *Semantics and pragmatics*, *2*, 4–1.
- [30] Goldberg, A. E. (2003). Constructions: a new theoretical approach to language. *Trends in cognitive sciences*, *7*, 219–224.
- [31] Golland, D., Liang, P., & Klein, D. (2010). A game-theoretic approach to generating spatial descriptions. In *Proceedings of the 2010 conference on empirical methods in natural language processing* (pp. 410–419). Association for Computational Linguistics.
- [32] Goodman, N. D., & Stuhlmüller, A. (2013). Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science*, *5*, 173–184.
- [33] Graf, C., Hawkins, R. X. D., & Goodman, N. D. (2016). Animal, dog, or dalmatian? Level of abstraction in nominal referring expressions. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*.
- [34] Grice, H. (1975). *Logic and conversation*.
- [35] Grodner, D. J., Klein, N. M., Carbary, K. M., & Tanenhaus, M. K. (2010). Some, and possibly all, scalar inferences are not delayed: Evidence for immediate pragmatic enrichment. *Cognition*, *116*, 42–55.
- [36] Heim, I., & Kratzer, A. (1998). *Semantics in generative grammar* volume 13. Blackwell Oxford.
- [37] Horn, L. R. (1984). Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature. In D. Schiffrin (Ed.), *Meaning, Form, and Use in Context: Linguistic Applications* (pp. 11–42). Washington, D.C.: Georgetown University Press.
- [38] Horowitz, A. C., & Frank, M. C. (2016). Children’s pragmatic inferences as a route for learning about the world. *Child Development*, *87*, 807–819.
- [39] Huang, Y. T., & Snedeker, J. (2009). Online interpretation of scalar quantifiers: Insight into the semantics–pragmatics interface. *Cognitive psychology*, *58*, 376–415.

- [40] Jackendoff, R. S. (2002). *Foundations of language: Brain, meaning, grammar, evolution*. Oxford, UK: Oxford University Press.
- [41] Jäger, G. (2008). Applications of game theory in linguistics. *Language and Linguistics compass*, 2, 406–421.
- [42] Jones, M., & Love, B. C. (2011). Bayesian fundamentalism or enlightenment? on the explanatory status and theoretical contributions of bayesian models of cognition. *Behavioral and Brain Sciences*, 34, 169–188.
- [43] Kao, J. T., Bergen, L., & Goodman, N. D. (2014). Formalizing the pragmatics of metaphor understanding. In *Proceedings of the 36th annual meeting of the Cognitive Science Society* (pp. 719–724).
- [44] Kao, J. T., & Goodman, N. D. (2015). Let’s talk (ironically) about the weather: Modeling verbal irony. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*.
- [45] Kao, J. T., Wu, J. Y., Bergen, L., & Goodman, N. D. (2014). Nonliteral understanding of number words. *Proceedings of the National Academy of Sciences*, 111, 12002–12007.
- [46] Kemp, C., & Regier, T. (2012). Kinship categories across languages reflect general communicative principles. *Science*, 336, 1049–1054.
- [47] Keysar, B., Lin, S., & Barr, D. J. (2003). Limits on theory of mind use in adults. *Cognition*, 89, 25–41.
- [48] Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences*, 105, 10681–10686.
- [49] Kirby, S., Tamariz, M., Cornish, H., & Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition*, 141, 87–102.
- [50] Lane, L. W., Groisman, M., & Ferreira, V. S. (2006). Don’t talk about pink elephants! speakers’ control over leaking private information during language production. *Psychological science*, 17, 273–277.
- [51] Lassiter, D., & Goodman, N. D. (2015). Adjectival vagueness in a bayesian model of interpretation. *Synthese*, (pp. 1–36).
- [52] Levelt, W. J. (1993). *Speaking: From intention to articulation* volume 1. MIT press.
- [53] Levinson, S. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. MIT Press.

- [54] Levy, R. (2008). A noisy-channel model of rational human sentence comprehension under uncertain input. In *Proceedings of the conference on empirical methods in natural language processing* (pp. 234–243). Association for Computational Linguistics.
- [55] Lewis, D. (1969). *Convention: A philosophical study*. John Wiley & Sons.
- [56] Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York, NY: Henry Holt and Co.
- [57] Monroe, W., & Potts, C. (2015). Learning in the Rational Speech Acts model. In *Amsterdam Colloquium*. Amsterdam: ILLC.
- [58] Nordmeyer, A. E., & Frank, M. C. (2014). A pragmatic account of the processing of negative sentences. In *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*.
- [59] Noveck, I. (2001). When children are more logical than adults: experimental investigations of scalar implicature. *Cognition*, 78, 165–188. doi:10.1016/S0010-0277(00)00114-1.
- [60] Noveck, I. A., & Reboul, A. (2008). Experimental pragmatics: A gricean turn in the study of language. *Trends in Cognitive Sciences*, 12, 425–431.
- [61] Orita, N., Vornov, E., Feldman, N. H., & Boyd-Graber, J. (2014). Quantifying the role of discourse topicality in speakers’ choices of referring expressions. *ACL 2014*, (p. 63).
- [62] Papafragou, A., & Musolino, J. (2003). Scalar implicatures: experiments at the semantics-pragmatics interface. *Cognition*, 86, 253–282. doi:10.1016/S0010-0277(02)00179-8.
- [63] Perfors, A., & Navarro, D. J. (2014). Language evolution can be shaped by the structure of the world. *Cognitive science*, 38, 775–793.
- [64] Piantadosi, S. T., Tily, H., & Gibson, E. (2012). The communicative function of ambiguity in language. *Cognition*, 122, 280–291.
- [65] Potts, C., Lassiter, D., Levy, R., & Frank, M. C. (2015). Embedded implicatures as pragmatic inferences under compositional lexical uncertainty. *Journal of Semantics*, (p. ffv012).
- [66] Qing, C., & Franke, M. (2015). Variations on a bayesian theme: Comparing bayesian models of referential reasoning. In *Bayesian natural language semantics and pragmatics* (pp. 201–220). Springer.
- [67] Regier, T., Kay, P., & Khetarpal, N. (2007). Color naming reflects optimal partitions of color space. *Proceedings of the National Academy of Sciences*, 104, 1436–1441.

- [68] Roberts, C. (1996). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Working Papers in Linguistics-Ohio State University Department of Linguistics*, (pp. 91–136).
- [69] Russell, B. (2012). *Probabilistic reasoning and the computation of scalar implicatures*. Ph.D. thesis Brown University.
- [70] Smith, N. J., Goodman, N., & Frank, M. (2013). Learning and using language via recursive pragmatic reasoning about other agents. In *Advances in Neural Information Processing Systems*. URL: <http://langcog.stanford.edu/papers/SGF-NIPS2013.pdf>.
- [71] Sperber, D., & Wilson, D. (1986). *Relevance: Communication and Cognition*. Oxford, UK: Blackwell Publishers.
- [72] Stiller, A., Goodman, N. D., & Frank, M. C. (2011). Ad-hoc scalar implicature in adults and children. In *Proceedings of the 33rd annual meeting of the Cognitive Science Society* (pp. 2134–2139). Cognitive Science Society Austin, TX.
- [73] Stiller, A. J., Goodman, N. D., & Frank, M. C. (2015). Ad-hoc implicature in preschool children. *Language Learning and Development*, 11, 176–190.
- [74] Tenenbaum, J., Kemp, C., Griffiths, T., & Goodman, N. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331, 1279.
- [75] Viethen, J., & Dale, R. (2006). Algorithms for generating referring expressions: do they do what people do? In *Proceedings of the Fourth International Natural Language Generation Conference* (pp. 63–70). Association for Computational Linguistics.
- [76] Vogel, A., Bodoia, M., Potts, C., & Jurafsky, D. (2013). Emergence of gricean maxims from multi-agent decision theory. In *HLT-NAACL* (pp. 1072–1081).
- [77] Vogel, A., Potts, C., & Jurafsky, D. (2013). Implicatures and nested beliefs in approximate decentralized-pomdps. In *ACL (2)* (pp. 74–80). Citeseer.
- [78] Vouloumanos, A., Onishi, K., & Pogue, A. (2012). Twelve-month-old infants recognize that speech can communicate unobservable intentions. *Proceedings of the National Academy of Sciences*, 109, 12933–12937.
- [79] Williamson, T. (2002). *Vagueness*. Routledge.
- [80] Xu, Y., & Regier, T. (2014). Numeral systems across languages support efficient communication: From approximate numerosity to recursion. *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*, .
- [81] Yoon, E. J., Tessler, M. H., Goodman, N. D., & Frank, M. C. (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*.