

Understanding the Rational Speech Act model

Arianna Yuan,¹ Will Monroe,² Yu Bai,³ and Nate Kushman⁴

Departments of ¹Psychology, ²Computer Science, ³Statistics, Stanford University, Stanford, CA 94305

⁴Microsoft Research, Cambridge, CB1 2FB, UK

Abstract

The Rational Speech Act (RSA) model, which proposes that probabilistic speakers and listeners recursively reason about each other’s mental states to communicate, has been successful in explaining many pragmatic reasoning phenomena. However, several theoretical questions remain unanswered. First, will such a pragmatic speaker–listener pair always outperform their literal counterparts who do not reason about each others mental states? Second, how does communication effectiveness change with the number of recursions? Third, when exact inference cannot be performed, how does limiting the computational resources of the speaker and listener affect these results? We systematically analyzed the RSA model and found that in Monte Carlo simulations pragmatic listeners and speakers always outperform their literal counterparts and the expected accuracy increases as the number of recursions increases. Furthermore, limiting the computation resources of the speaker and listener so they sample only the top k most likely options leads to *higher* expected accuracy. We verified these results on a previously collected natural language dataset in color reference games. The current work supplements the existing RSA literature and could guide future modeling work.

Keywords: Pragmatic Reasoning. Rational Speech Act model.

Introduction

The Rational Speech Act (RSA) model has successfully explained a number of psycholinguistic findings in pragmatic reasoning (Goodman & Stuhlmüller, 2013; Goodman & Frank, 2016; Frank, 2016; Kao, Bergen, & Goodman, 2014). This model proposes that probabilistic speakers and probabilistic listeners recursively reason about each other’s mental states to infer the meaning of utterances and generate utterances in response. One simple example demonstrating how RSA works is a reference game scenario (Frank & Goodman, 2012) in which the speaker and the listener can both see three faces (Figure 1): one with hat and glasses (HG), one with only glasses (G), and one with neither (N); one of these is the speaker’s “friend”. The speaker says “my friend has glasses”, presupposing that there is a single friend. The listeners, who know that the only alternative utterance was “my friend has a hat”, share the intuition that the sentence “my friend has glasses” refers to G and not HG or N (Stiller, Goodman, & Frank, 2011). The RSA model provides a nice explanation for that intuition, as it proposes that the listener reasons about the speaker’s mental state, and realizes that if the speaker meant to refer to the one with hat and glasses (HG), he/she could have said “my friend has a hat” to avoid any ambiguity. The fact that he/she did not say so implied that HG was not the referent.

Despite the success of the RSA model in explaining a variety of pragmatic reasoning phenomena, several

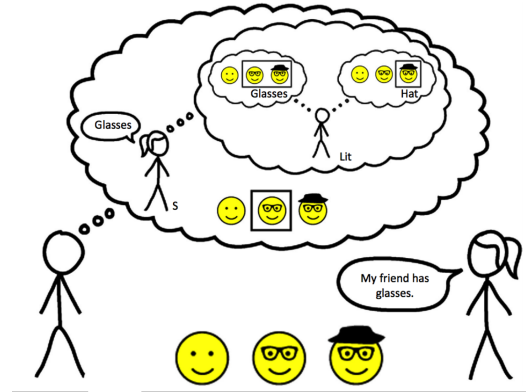


Figure 1: Pragmatic reasoning in a reference game. The speaker says my friend has glasses. The listener needs to guess which one is the speaker’s friend. Adapted from Goodman and Frank (2016).

theoretical questions remain unanswered. In this paper we address three relevant questions:

(1) **Will pragmatic listeners and pragmatic speakers always communicate better than literal listeners and literal speakers? If so, what is the reason underlying their advantage?** Despite broad empirical evidence showing that people tend to behave more like pragmatic listeners and speakers, it’s not clear whether we should always expect that recursive reasoning will provide an improvement.

(2) **How does the number of recursions affect the accuracy of agents using the RSA model?** In previous modeling work using the RSA model, the depth of recursions was usually set to two (but see Bergen, Goodman, and Levy (2012) and Degen, Franke, and Jager (2013)). However, if we expect a general advantage of applying pragmatic reasoning in communication, we should expect more recursions to result in even higher performance in reference games. Although some studies have shown that deep recursion may not be realistic in reference games (Degen & Franke, 2012), more recent evidence has shown deeper recursion in some participants (Franke & Degen, 2016), consistent with work on economic games where higher recursion depth in human Theory of Mind is sometimes found (Camerer, Ho, & Chong, 2004). In the current paper we extend previous work that has studied deeper recursion on small meaning matrices to large and realistic meaning matrices (Franke & Degen, 2016).

(3) **How do computational constraints on exact calculation of the probabilities in the RSA model affect its performance?** Prior work on RSA has focused on the

setting where the context set is small, such that it is possible to normalize the posterior probability over all items in the candidate set. However, if we include all possible world states in the candidate set for the listener model and all possible utterances in the candidate set for the speaker model, it would be impossible, or at least unrealistic, to operate on the entire candidate set when performing normalization. Therefore, we explore the effect on accuracy of using only the k most likely candidates when performing the calculations at each step.

In this paper, we first briefly review the formalization of a classic RSA model. Then we show that we can simplify the model under certain assumptions and derive a concise expression of the expected accuracy in a reference game given any listener and speaker models. Using this formula, we run Monte Carlo simulations on random meaning functions and present our results that address the three questions mentioned above. The key finding is that as listeners and speakers perform more recursions of inferring each other’s mental states, their joint performance gets better. We also verify that these conclusions hold on a dataset of human utterances collected in a color reference game (Monroe, Hawkins, Goodman, & Potts, 2017). Finally, we discuss the implications of our findings.

Models

A classic RSA model (Goodman & Frank, 2016) usually starts from a literal listener:

$$L_0(t|u, M) \propto M(u, t)P(t), \quad (1)$$

where t is a possible world in a set of all possible worlds C , u is an utterance drawn from a set of possible utterances U , P is a prior over worlds, and $M = M(u, t)$ is a meaning function that takes the value 1 if u is true of t , otherwise 0.

A pragmatic speaker would infer a literal listener model from the meaning function M and then build his/her own model accordingly:

$$S_1(u|t, M) \propto e^{\alpha(\log(L_0(t|u, M)) - \kappa(u))}. \quad (2)$$

Here, $\kappa(u)$ is a real-valued cost function on utterances, and $\alpha \in [0, \infty)$ is an inverse temperature parameter controlling how rational the speaker is. Specifically, if α is large, the speaker will choose the utterance with highest likelihood, whereas when α is small, the speaker tends to choose utterances more randomly and thus suboptimally. A pragmatic listener in turn builds his/her own listener model L_2 based on the pragmatic speaker model S_1 :

$$L_2(t|u, M) \propto S_1(u|t, M)P(t) \quad (3)$$

For the sake of simplicity, we assume the prior over worlds is uniform, so equation (1) becomes:

$$L_0(t|u, M) \propto M(u, t). \quad (4)$$

We also assume the cost function $\kappa(u)$ is a constant function and $\alpha = 1$, which amounts to a scenario in which the speaker and the listener use a matching-probability strategy.

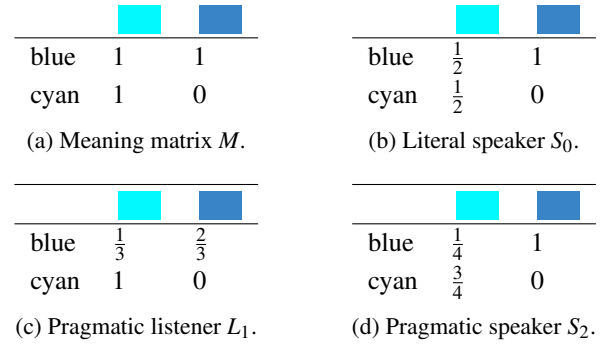


Figure 2: Simulating pragmatic reasoning in a reference game with a basic RSA model.

Under these assumptions, we can simplify the models to:

$$L_0(t|u, M) \propto M(u, t) \quad (5)$$

$$S_1(u|t, M) \propto L_0(t|u, M) \quad (6)$$

$$L_2(t|u, M) \propto S_1(u|t, M) \quad (7)$$

Similarly, a speaker model can also be derived from the meaning function:

$$S_0(u|t, M) \propto M(u, t) \quad (8)$$

$$L_1(t|u, M) \propto S_0(u|t, M) \quad (9)$$

$$S_2(t|u, M) \propto L_1(t|u, M) \quad (10)$$

In effect, each step in the recursion simply normalizes the M across the world states, or across the utterances. To illustrate how this simulates pragmatic reasoning, we present a concrete example. In the meaning matrix in Figure 2a, we see that the word “blue” applies to both of the two colors and the word “cyan” only applies to the left one. Therefore, a literal speaker would be equally likely to use “blue” and “cyan” to refer to the left color. Formally, we can normalize the columns of the meaning matrix (Equation 8) to get the literal speaker model S_0 (Figure 2b). A pragmatic listener L_1 builds a mental representation of the literal speaker model S_0 and uses Bayes’ rule to interpret the speaker’s utterance. For example, if the listener hears the word “blue”, the chance that “blue” is produced for the left color is only half of the chance that it is produced for the right color. Therefore, a pragmatic listener would assign a probability of $\frac{1}{3}$ to the left color and $\frac{2}{3}$ to the right color. Formally, the pragmatic listener normalizes the rows of the literal speaker model S_0 to obtain its own model L_1 (Figure 2c). Finally, a pragmatic speaker mentally simulates the pragmatic listener L_1 and chooses utterances accordingly, i.e., it normalizes the columns of L_1 (Equation 10) to get its own model S_2 (Figure 2d).

As shown in Figure 2, the listener model and the speaker model change after each iteration. For a 2x2 meaning matrix, the ultimate converged matrix can be explicitly derived as:

$$M = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \xrightarrow{n \rightarrow \infty} \begin{bmatrix} p & 1-p \\ 1-p & p \end{bmatrix}.$$

	0	1
blue	0	1
cyan	1	0

Figure 3: The converged listener model and speaker model after many recursions using the RSA model.

when $ad \neq 0$ or $bc \neq 0$.

This is because $\lim_{n \rightarrow \infty} L_n = \lim_{n \rightarrow \infty} S_n$ and both the rows and the columns sum up to 1. To explicitly derive p as a function of a, b, c, d , it can be shown that the ratio $\frac{ad}{bc}$ (or $\frac{bc}{ad}$ if $bc = 0$) remains unchanged under row and column normalizations, so we have

$$\frac{p^2}{(1-p)^2} = \frac{ad}{bc},$$

thus

$$p = \frac{\sqrt{ad}}{\sqrt{ad} + \sqrt{bc}} \quad (11)$$

Using this formula, we know that after many iterations, the pragmatic listener and speaker in Figure 2 will both converge to the matrix in Figure 3, resulting in a one-to-one mapping between utterances and worlds, i.e., a better communication protocol. This implies that the recursions in the RSA model discourage distributions containing non-specific utterances.

Does the observation that RSA optimizes the performance of a listener-speaker pair in the 2x2 case generalize to n -dimensional meaning matrices? To answer this question, we run some simulations using both random and naturalistic meaning matrices. We first give a mathematical formulation of the problem. Starting from a meaning function, represented by a binary matrix $M \in \mathbb{R}^{p \times q}$, we define two operators. The first is “row normalization”, $\phi_L \in \mathbb{R}^{p \times q} \rightarrow \mathbb{R}^{p \times q}$, which is used to derive the listener model L . The second is “column normalization” $\phi_S \in \mathbb{R}^{p \times q} \rightarrow \mathbb{R}^{p \times q}$, which is used to derive the speaker model S . According to the basic RSA model, which assumes the prior over worlds is uniform and the cost function of each utterance is constant, we have:

$$\begin{aligned} L_0 &= \phi_L(M) \\ S_0 &= \phi_S(M) \\ L_n &= \phi_L(S_{n-1}) \\ S_n &= \phi_S(L_{n-1}) \end{aligned} \quad (12)$$

This alternating normalization is also known as the Sinkhorn-Knopp algorithm (Sinkhorn & Knopp, 1967), which is shown to converge if and only if A has a positive diagonal, i.e. there exists a permutation σ of $\{1, \dots, n\}$ such that $A_{i, \sigma_i} > 0$ for all i . This is a very reasonable assumption for a meaning matrix, since it implies that every world state has at least one utterance which refers to it and every utterance refers to at least one possible world state.

Under this framework, we can derive the expected accuracy of a listener in a reference game given any meaning matrix M

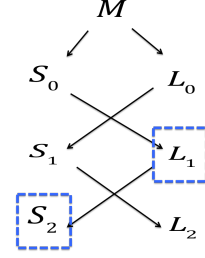


Figure 4: Evaluating the joint performance by a listener model L_{n-1} and a speaker model S_n , where n is the number of recursions.

and level of recursion n . This derivation allows us to examine whether the joint performance by a listener-speaker pair is better at higher levels n . We simulate the reference game in the following way: 1. Sample a true world w from a uniform distribution. 2. Given w , sample an utterance u from the speaker S . 3. Given u , sample a world prediction \hat{w} from the listener L . The accuracy of the pair is the probability that $\hat{w} = w$, i.e., the listener model makes a correct response.

Formally, given a world w_j , the speaker S picks an utterance u , according to the probability distribution $S_n e_j$, where $e_j = [0 \dots 1 \dots 0]^T \in \mathbb{R}^q$, and e_j 's j th entry is 1. As illustrated by the blue squares with dashed borders in Figure 4, we pair up a listener with its derived pragmatic speaker and compute their joint performance. In other words, we calculate the performance of pair $\langle S_n, L_{n-1} \rangle$. So the listener L guesses a t according to the listener model given the utterance the speaker used: $L_{n-1}^T S_n e_j \in \mathbb{R}^q$. The accuracy is the probability that the listener guesses the world w_j , i.e.,

$$[L_{n-1}^T S_n e_j]_j = e_j^T L_{n-1}^T S_n e_j = (L_{n-1}^T S_n)_{jj}$$

As the prior over the world w_1, \dots, w_q is uniform, the expected accuracy, a , can be written as:

$$a = \frac{1}{q} ((L_{n-1}^T S_n)_{11} + \dots + (L_{n-1}^T S_n)_{qq}) = \frac{1}{q} \text{tr}(L_{n-1}^T S_n)$$

Using this formula, we first explore how the number of recursions affects the expected accuracy.

Simulations on Random Matrices

We first present the results of simulations using random meaning functions. We generate meaning matrices $\{0, 1\}^{p \times p}$ for $p = \{10, 20, 30, 40, 50\}$ and each entry independently containing 1 with probability 0.1. For each matrix size,¹ we randomly generate 200 matrices and compute the expected accuracy of $\langle S_n, L_{n-1} \rangle$, for $n = 1, 2, \dots, 50$. The first thing we find is that regardless of the size of the meaning matrices,

¹The number of rows and the number of columns do not have to be equal. Here we just use square matrices for simplicity, and the results are qualitatively the same for non-square matrices.

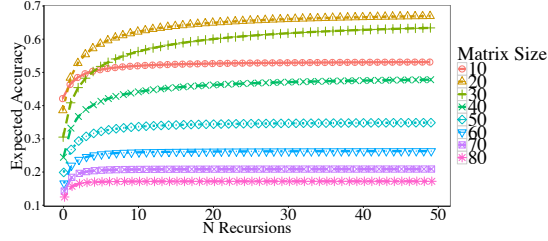


Figure 5: Joint reference game accuracy of $\langle L_{n-1}, S_n \rangle$.

more recursions always results in higher performance (Figure 5). Secondly, although we simulate up to 50 recursions, the gain of doing recursive pragmatic reasoning is only salient in the first few recursions, especially when the meaning matrix size is large. This is promising for cognitive plausibility, since we don't expect people to perform very many recursions (due to limited working memory capacity) and it is important that the benefit of pragmatic reasoning emerges with a small number of recursions.

Effect of Sampling Next we examine how sampling affects expected accuracy. When the speaker has a large vocabulary, it is no longer plausible to enumerate all the utterances and do column-wise normalization. Instead, he/she may select only the top k most likely utterances and discard the rest for further computation, i.e., set the probability of the others to be 0. The subsequent normalization is then applied only to the vector of the remaining non-zero probabilities. More formally, for any column $a \in \mathbb{R}^{q \times 1}$, $a_i \geq 0$, we convert a_i to \tilde{a}_i , where:

$$\tilde{a}_i = \frac{a_i \mathbb{I}\{a_i \text{ is among the top } k \text{ elements in } a\}}{\sum_j a_j \mathbb{I}\{a_j \text{ is among the top } k \text{ elements in } a\}}.$$

Similarly, the listener model can consider only the top k most likely worlds and will set the probability of the rest of the worlds to be zero. This process approximates the expectation of another sampling-based approach in work on approximate Bayesian cognition (Goodman, Tenenbaum, Feldman, & Griffiths, 2008), i.e., agents sample a small set of k elements from the full distribution to realize.

We examine our top- k normalization procedure in two conditions, "speaker only" and "both". In the "speaker only" condition, only the speaker selects the top k utterances and perform truncated column-wise normalization, whereas in the "both" condition, both the speaker and the listener model select the top k elements in their probability distribution and perform this truncated normalization.

The results are shown in Figure 6. We find that in the "both" condition, as k decreases, the expected accuracy increases dramatically. When $k = 1, 2$, with only the first few iterations the accuracy rises rapidly. This is because forcing both the listener model and the speaker model to choose only the top few options allows them to quickly form a one-to-one mapping between utterances and worlds. We also found that, in general, having truncated normalization in both

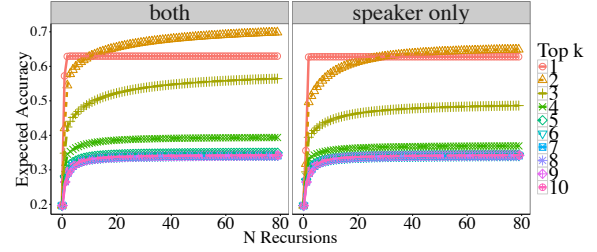


Figure 6: Expected accuracy of listener-speaker dyads when both agents (left) or only the speaker (right) consider the k most likely candidates.

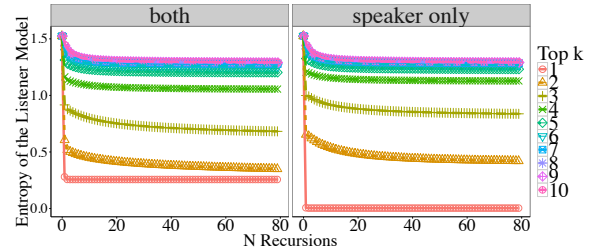


Figure 7: Entropy of listener-speaker dyads when both of them consider only the k most likely candidates (left) or only the speaker considers the k most likely candidates (right).

the listener model and the speaker model results in higher joint performance. At first this may seem counter-intuitive, since we naturally expect a larger sampling size to lead to better reasoning. However, as shown here, when both the speaker and the listener are aware of each other's cognitive constraints and the desire to keep only a small number of possibilities in their individual mental model, it actually helps the dyad to quickly converge to a convention that approximates one-to-one (or few-to-few) mapping.

Supporting this, we found that more recursions reduces the mean entropy of both the speaker model (Figure 7) and the listener model (not plotted here). This explains the higher accuracy associated with more recursions: at higher levels, listeners and speakers reduce their uncertainty in each other's world-utterance mappings.

An interesting finding is that $k = 2$ leads to even better accuracy than $k = 1$. To understand the underlying mechanism, we examine how the number of non-zero elements for the listener and speaker changes as a function of k and number of recursions (Figure 8). We find that, not surprisingly, choosing only the top 1 option quickly reduces the number of non-zero elements to 50, which means the speaker and the listener both form a one-to-one mapping between worlds and utterances. However, it is likely that their mappings are different, which will lower their joint performance. Note that this top- k selection process is irreversible, in the sense that once an element is set to zero, it won't become non-zero in later recursions. That explains why top-1 accuracy is lower than the top-2 accuracy, as the

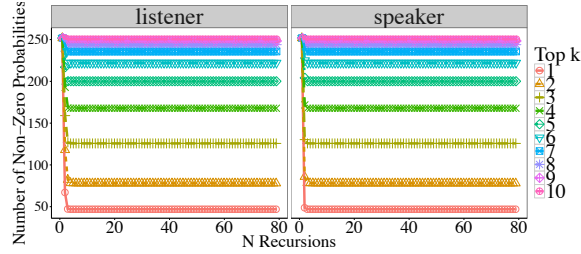


Figure 8: Number of non-zero elements in the listener model L_i (left) and the speaker model S_j (right) as a function of sampling size k and number of recursions n .

top-2 sampling process allows more elements to be non-zero at each iteration, thus allowing the speaker and listener to coordinate through more recursions and jointly optimize their models. Another way to interpret this result is that $k = 2$ may strike a good balance between exploration and exploitation, as found in previous game-theoretic work (Franke & Jäger, 2014). When k is even larger, the benefit of being flexible is canceled out by the cost of being less specific, as seen in the low expected accuracies of models with large k (Figure 6).

Simulations on a Naturalistic Dataset

In the previous section we ran simulations on binary random matrices. Next, we verify our results on a realistic meaning matrix extracted from data collected in a color reference game by Monroe et al. (2017). They paired Amazon Mechanical Turkers into dyads, and one of the Turkers was assigned the role of a speaker and the other was assigned a role of a listener. The speaker could see a target color alongside two distractors and needed to describe the color to the listener so that the listener could identify the target correctly. We limit our investigation to single-word utterances that appear at least twice in this corpus, a total of 261 utterances. We also divide colors on a continuous spectrum into 128 color categories. We count the frequency of each unigram-color pair in the corpus and obtained the co-occurrence matrix. We treat this as our meaning matrix $M \in \mathbb{R}^{261 \times 128}$. Note that this meaning matrix is no longer binary, but integer-valued, containing the information of the utterance prior and the world prior.²

We ran the same simulations on this realistic meaning matrix and the results are shown in Figure 9. Consistent with what we found in the simulations on random matrices, we can see that the accuracy increases as the parameter k decreases. Limiting k to be small (e.g., 1 or 2) results in a communication protocol that achieves high accuracy.

A possible concern is that while more recursions allow the listener and speaker to form a better communication protocol, they may end up with a world-utterance mapping that is only interpretable by themselves, and not sensible to other people. In other words, more recursions may result in a very efficient

²Strictly, this is not a literal meaning matrix, but reflects pragmatic use. However, treating it as a pragmatic speaker model (e.g., S_1) does not alter the qualitative results of the simulation.

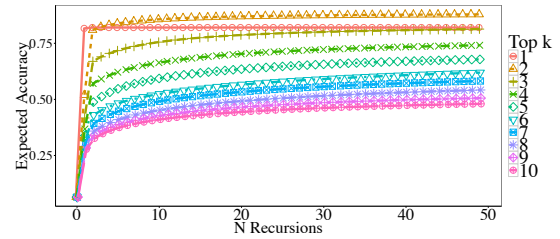


Figure 9: Expected accuracy calculated from a meaning function, which is extracted from human data in a color reference game (Monroe et al., 2017).

cyan	0.03	0	0	0	0.01
blue-green	0.02	0.01	0	0	0.01
blue-grey	0	0	0.01	0	0
blue-purple	0	0	0	0.01	0
bluish	0	0	0.01	0.01	0.02

(a) A submatrix of the original probabilistic speaker model S_0 .

cyan	0.21	0	0	0	0
blue-green	0.08	0.26	0	0	0.03
blue-grey	0	0	0.53	0	0
blue-purple	0	0	0	0.27	0
bluish	0	0	0	0	0.36

(b) A submatrix of the converged probabilistic speaker model S_n .

Figure 10: Examples showing how probabilities in the speaker model change from the original values to the ultimate values in the converged matrix.

communication protocol that only works for this specific pair. We argue that this won't be the case, since the top k truncated normalization procedure will guarantee that the support of L_n is a subset of the support of S_{n-1} , i.e., recursive pragmatic reasoning does not create new world-utterance associations, but only changes the relative strength of existing associations.

In Figure 10, we show a submatrix of the original speaker model and the corresponding submatrix of the converged speaker model. Note that the normalization was still performed over the entire matrix (all utterances and all worlds). We can see that the original speaker model S_0 becomes more specific after it converges to S_n . In particular, in the original speaker model S_0 , the non-zero probabilities are very small, indicating that the distribution is spread over many utterances. However, after a few iterations using the RSA model, some utterances are assigned large probabilities and thus dominate the distribution.

Discussion

In this work, we systematically evaluate how expected accuracy of a pragmatic listener-speaker dyad in a reference game changes as a function of: (a) size of vocabulary and world state space, (b) number of pragmatic reasoning recursions and (c) number of candidates considered at each iteration. We find that regardless of vocabulary size and number of possible worlds, more pragmatic reasoning recursions always lead to higher accuracy. This is explained by reduction in entropy, and it is consistent with previous theoretical work on game-theoretical pragmatics, which has shown that for agents that maximize expected utility, in each step of recursive iterated pragmatic reasoning, expected communicative success is non-decreasing (Jäger, 2011). Our work confirms this point in the soft-utility-maximization setting. In addition, we find that considering only a few candidates at each iteration in the recursion helps a listener-speaker dyad play the reference task better.

Our work is closely related to multi-agent communication, a topic that has recently received significant attention in the machine learning community (Andreas, Dragan, & Klein, 2017; Havrylov & Titov, 2017). For instance, several previous efforts have focused on training a speaker agent that can describe a target image in natural or synthetic language, or to enable a listener agent to identify the target image (Andreas & Klein, 2016; Lazaridou, Peysakhovich, & Baroni, 2017). Andreas and Klein trained a speaker model to generate a context-sensitive caption for a target image in a reference game. However, the speaker does not explicitly model the listener model. In other words, their model used only one round of recursion. According to our simulation, more recursions is likely to yield better communication protocols. Therefore, future work should explore whether adding more recursions in such neural pragmatic speaker and listener models would lead to performance improvements.

One of the limitations of the current work is that we assume the speaker and the listener start from exactly the same meaning matrix, which is a common practice in previous RSA modeling work. However, in reality, a listener and speaker might start from slightly different meaning functions. Future work should examine how small mismatches between the listener and speaker's meaning functions might change the conclusions derived in the current work.

References

- Andreas, J., Dragan, A., & Klein, D. (2017). Translating neuralese. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics* (pp. 232–242).
- Andreas, J., & Klein, D. (2016). Reasoning about pragmatics with neural listeners and speakers. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing* (pp. 1173–1182).
- Bergen, L., Goodman, N., & Levy, R. (2012). That's what she (could have) said: How alternative utterances affect language use. In *Proceedings of the Annual Meeting of the Cognitive Science Society*.
- Camerer, C. F., Ho, T.-H., & Chong, J.-K. (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3), 861–898.
- Degen, J., & Franke, M. (2012). Optimal reasoning about referential expressions. *Proceedings of SemDIAL*.
- Degen, J., Franke, M., & Jäger, G. (2013). Cost-based pragmatic inference about referential expressions. In *Proceedings of the Annual Meeting of the Cognitive Science Society*.
- Frank, M. C. (2016). Rational speech act models of pragmatic reasoning in reference games. *PsyArXiv*.
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998–998.
- Franke, M., & Degen, J. (2016). Reasoning in reference games: Individual-vs. population-level probabilistic modeling. *PLoS one*, 11(5), e0154854.
- Franke, M., & Jäger, G. (2014). Pragmatic back-and-forth reasoning. In *Pragmatics, semantics and the case of scalar implicatures* (pp. 170–200). Springer.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11), 818–829.
- Goodman, N. D., & Stuhlmüller, A. (2013). Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science*, 5(1), 173–184.
- Goodman, N. D., Tenenbaum, J. B., Feldman, J., & Griffiths, T. L. (2008). A rational analysis of rule-based concept learning. *Cognitive science*, 32(1), 108–154.
- Havrylov, S., & Titov, I. (2017). Emergence of language with multi-agent games: learning to communicate with sequences of symbols. In *Advances in Neural Information Processing Systems* (pp. 2146–2156).
- Jäger, G. (2011). Game-theoretical pragmatics. *Handbook of logic and language*, 467–491.
- Kao, J., Bergen, L., & Goodman, N. (2014). Formalizing the pragmatics of metaphor understanding. In *Proceedings of the Annual Meeting of the Cognitive Science Society*.
- Lazaridou, A., Peysakhovich, A., & Baroni, M. (2017). Multi-agent cooperation and the emergence of (natural) language. In *Proceedings of the International Conference on Learning Representations*.
- Monroe, W., Hawkins, R. X., Goodman, N. D., & Potts, C. (2017). Colors in context: A pragmatic neural model for grounded language understanding. *Transactions of the Association of Computational Linguistics*, 5(1), 325–338.
- Sinkhorn, R., & Knopp, P. (1967). Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21(2), 343–348.
- Stiller, A., Goodman, N., & Frank, M. (2011). Ad-hoc scalar implicature in adults and children. In *Proceedings of the Annual Meeting of the Cognitive Science Society*.