

A Rate–Distortion view of human pragmatic reasoning

Noga Zaslavsky^{1,2}

Jennifer Hu¹

Roger P. Levy¹

¹Department of Brain and Cognitive Sciences, ²Center for Brains Minds and Machines

Massachusetts Institute of Technology

{nogazs, jennhu, rplevy}@mit.edu

What computational principles underlie human pragmatic reasoning? A prominent approach to pragmatics is the Rational Speech Act (RSA) framework (Frank and Goodman, 2012; Goodman and Frank, 2016), which formulates pragmatic reasoning as probabilistic speakers and listeners recursively reasoning about each other with the goal of cooperatively gaining communicative utility. While RSA enjoys broad empirical support, much remains unknown about the dynamics of RSA recursion and whether it can be characterized by a general optimization principle. It has been conjectured that RSA dynamics is guaranteed to increase expected utility (e.g., Yuan et al., 2018; Peloquin et al., 2019), but these explorations have relied on numeric simulations, leaving open key questions about the dynamics of RSA models.

In this work we present a set of analytic results, demonstrated by model simulations, that extend the mathematical understanding of the RSA framework and ground it in Rate–Distortion (RD) theory (Shannon, 1948). First, we show that RSA recursion optimizes a tradeoff between expected utility and communicative effort, disconfirming the conjecture that expected utility is guaranteed to improve with recursion depth. Second, we show that RSA can be grounded in RD theory, while maintaining a similar ability to account for human behavior and avoiding a bias of RSA toward random utterance production. Taken together, these results suggest that human pragmatic reasoning may be understood in terms of RD theory.

RSA as Alternating–Maximization. In RSA, the speaker is defined by a production distribution $S(u|m)$ over possible utterances u given meaning m , and the listener is defined by an inference distribution $L(m|u)$. RSA recursively relates the speaker and listener (see Figure 1) by assuming a Bayesian listener— $L(m|u) \propto S(u|m)P(m)$,

with $P(m)$ a prior distribution on speaker meanings that is assumed to be in common ground—and a speaker that is bounded-rational with respect to a utility function $V(u, m)$, typically defined as $V(m, u) = \log L(m|u) - C(u)$ where $C(u)$ specifies the cost of u . That is, $S(u|m) \propto \exp(\alpha V(u, m))$, where α controls the degree to which the speaker maximizes utility.

Our first theoretical result is that RSA’s recursive reasoning implements an alternating maximization (AM) algorithm (Csiszár and Shields, 2004). However, this optimization does not maximize expected utility as previously conjectured, but rather a tradeoff between maximizing expected utility, $\mathbb{E}_S[V_L]$, and minimizing communicative effort measured by the conditional entropy of the speaker’s production distribution, $H_S(U|M)$, such that low effort corresponds to high entropy. Formally, we prove that for any $\alpha \geq 0$, each update step in RSA maximizes

$$\mathcal{G}_\alpha[S, L] = H_S(U|M) + \alpha \mathbb{E}_S[V_L]. \quad (1)$$

This analytic result is demonstrated numerically in Figure 1. Note that this does not imply that




Figure 1: Model simulations demonstrate that the RSA recursion implements an alternating maximization algorithm. RSA’s tradeoff \mathcal{G}_α improves with each speaker (blue) and listener (orange) update. Inset: Illustration of the RSA recursion.

the expected utility is necessarily maximized, and our analysis finds counter-examples where the expected utility decreases with recursion depth while the conditional entropy increases.

RD-RSA: Grounding RSA in Rate-Distortion.

In the communication setup of RSA, the speaker can be seen as a probabilistic encoder and the listener as a probabilistic decoder. From an information-theoretic perspective, RD theory predicts that the speaker and listener should jointly optimize the tradeoff between maximizing the expected utility and minimizing the number of bits required for communication. The latter is captured by the mutual information between speaker meanings and utterances, $I_S(M; U)$. Formally, this tradeoff is given by

$$\mathcal{F}_\alpha[S, L] = I_S(M; U) - \alpha \mathbb{E}_S[V_L], \quad (2)$$

which is closely related to \mathcal{G}_α and can similarly be optimized via an AM algorithm. The optimal RD-RSA listener is Bayesian, as in RSA; however, the optimal speaker takes the form $S(u|m) \propto S(u) \exp(\alpha V(u, m))$, differing from the RSA speaker in weighting the soft-max utility term by marginal utterance probability $S(u)$ (note that $S(u)$ is not pre-determined but rather changes with each iteration as the speaker reasons about the listener). We refer to this modified model of pragmatic reasoning as RD-RSA.

While RD-RSA and RSA are closely related, their theoretical motivation and precise predictions

differ. This raises the question: which model may better characterize human pragmatic reasoning? To begin to address this question, we study the dynamics of RSA and RD-RSA and compare their predictions against human behavior in pragmatic reference games (Vogel et al., 2014). Figure 2 shows simulated trajectories of pragmatic reasoning from the two models, and reveals that RD-RSA can account for human behavior as well as RSA. At the same time, our theoretical analysis reveals that the RSA speaker has an inherent bias toward non-informative utterance production, while the RD-RSA speaker does not.

Conclusions. We have shown that with a small adjustment, the RSA framework can be grounded in Shannon’s RD theory, while maintaining a similar ability to account for human data and avoiding a bias toward random utterance production. This work furthers the mathematical understanding of RSA models, and suggests that fundamental information-theoretic principles may give rise to human pragmatic reasoning.

References

I. Csiszár and P.C. Shields. 2004. Information theory and statistics: A tutorial. *Foundations and Trends in Communications and Information Theory*, 1(4):417–528.

Michael C Frank and Noah D Goodman. 2012. Predicting pragmatic reasoning in language games. *Science*, 336(6084):998–998.

Noah D Goodman and Michael C Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends Cogn. Sci.*, 20(11):818–829.

Benjamin N. Peloquin, Noah D. Goodman, and Michael C. Frank. 2019. The interactions of rational, pragmatic agents lead to efficient language structure and use. In *Proceedings of the 41st Annual Meeting of the Cognitive Science Society*, pages 912–917.

Claude E. Shannon. 1948. A mathematical theory of communication. *Bell System Technical Journal*, 27.

Adam Vogel, Andrés Gómez Emilsson, Michael C. Frank, Daniel Jurafsky, and Christopher Potts. 2014. Learning to reason pragmatically with cognitive limitations. In *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*, pages 3055–3060.

Arianna Yuan, Will Monroe, Yu Bai, and Nate Kushman. 2018. Understanding the Rational Speech Act model. In *Proceedings of the 40th Annual Conference of the Cognitive Science Society*, pages 2759–2764.

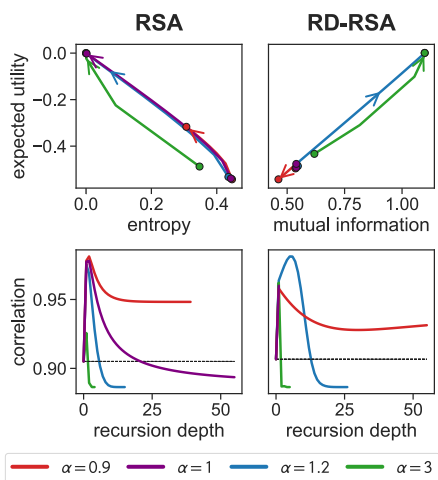


Figure 2: Top: Simulated model trajectories in RSA (left) and RD-RSA (right). Bottom: Pearson correlation between model predictions and human behavioral results from Vogel et al. (2014).