Title
The evolution of efficient compression in signaling games

Permalink
https://escholarship.org/uc/item/5dr5h4q0

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Author
Imel, Nathaniel

Publication Date
2023

Peer reviewed

The evolution of efficient compression in signaling games

Nathaniel Imel (nimel@uci.edu)
Department of Logic and Philosophy of Science
University of California, Irvine

Abstract

Converging evidence suggests that natural language meaning systems are *efficient* by jointly maximizing cognitive simplicity and communicative informativeness. Comparatively less is known about how languages might *optimize* over time for communicative efficiency. Our goal in this paper is to use minimal dynamic models to give a high-level description of the evolution of efficient meaning systems. To do this, we provide a model of emergent communication combining evolutionary game theory with a recent information theoretic account of efficiency in semantic systems. We perform simulations of adaptive dynamics requiring minimal assumptions about agents' cognitive resources, and observe that emergent languages converge near the achievable bounds of efficient compression. This unifies existing accounts of communicative efficiency with minimalist accounts of how meaning can emerge *ex nihilo*.

Keywords: simplicity/informativeness trade-off; information theory; language evolution; signaling games

Human languages use words to categorize the environment. Converging evidence suggests that our semantic category systems do this *efficiently* by maximizing simplicity and informativeness (Kemp & Regier, 2012; Kemp et al., 2018; Levinson, 2012). An important open question concerns how this optimization happens over time: how languages *optimize* for efficiency. Cognitive scientists have attended to this question using a variety of modeling frameworks, including cultural evolution (Carr et al., 2020; Carstensen et al., 2014; Kirby et al., 2015), semantic chaining (Xu et al., 2016), machine learning (Carlsson et al., 2021; Chaabouni et al., 2021; Kågebäck et al., 2020), and annealing processes (Zaslavsky et al., 2019; Zaslavsky et al., 2018). This paper aims to extend this progress using a simple but focused analysis. We ask: what is the relationship between communicative efficiency of semantic systems and their evolution in signaling games?

To answer this question, we model communication in signaling games and measure the information-theoretic efficiency of agents' semantic systems. Signaling games are evolutionary game-theoretic models widely applied across economics, biology, philosophy and linguistics (S. Huttegger et al., 2014; Nowak & Krakauer, 1999; Skyrms, 2010). In such games, a Sender and a Receiver coordinate via signals in order to maximize a joint payoff. Researchers often take the changes in players' strategies over repeated plays of the game to model of the emergence of communication. We focus on sim-max games, in which payoffs reflect the similarity of Sender's intended meaning to Receiver's reconstruction of that meaning (Franke & Correia, 2018; Jäger, 2007; Jäger & van Rooij, 2007; Komarova et al., 2007; O'Connor, 2014).

To explore the evolution of efficiency, we perform numerical simulations analyzing the optimality of communication systems resulting from two simple mechanisms for evolution in repeated games: Roth-Erev reinforcement learning and a discrete-time replicator dynamic (Erev & Roth, 1998; Hofbauer & Sigmund, 1998). To quantify communicative efficiency, we draw on a recent information-theoretic framework predicting efficient compression in natural language semantic systems (Zaslavsky, 2020; Zaslavsky et al., 2018). We measure the optimality of emergent communication systems in terms of their closeness to the Rate-Distortion curve.

We catalogue two empirical results.[1] First, signaling agents typically converge to languages that are near-optimal compression systems. This echoes recent results from emergent communication (Chaabouni et al., 2021; Tucker et al., 2022a, 2022b) that deep reinforcement learning in multi-agent reference games can lead to semantic systems that achieve efficient compression. Second, the degree to which perceptual distinctions are rewarded in the game typically constrains where, along the Rate-Distortion curve, languages will converge. This mirrors the findings of Chaabouni et al. (2021), lending further support to the idea that diverse solutions to the simplicity/informativeness trade-off may be accounted for in terms of diverse environmental need to discriminate stimuli. Our results are novel in indicating that some meaning systems that achieve efficient compression are also those yielded by well-studied evolutionary dynamics.

The Simplicity/Informativeness Trade-Off

Researchers have argued that meanings expressed in lexicons across the world are optimized for efficient communication. This idea can be summarized roughly as follows: a language can be simple (by e.g. containing a single expression) and a language can be informative (by e.g. having unique expressions for every posible thought). A language cannot, however, be both simple and informative: these two pressures trade off against each other. A hypothesis in linguistics is that the natural languages are near solutions to this multi-objective optimization problem, and that these efficiency pressures explain constraints on crosslinguistic semantic variation (Kemp

[1] Code for reproduction is available at https://github.com/nathimel/rdsg.

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society*.

& Regier, 2012; Kemp et al., 2018; Levinson, 2012). This efficient communication hypothesis has been successfully applied across semantic domains, including kinship terms, color terms, number terms, container terms, quantifiers, tense and evidentiality, boolean connectives, person systems, indefinite pronouns and modals (Denić et al., 2022; Imel & Steinert-Threlkeld, 2022; Kemp & Regier, 2012; Mollica et al., 2021; Regier et al., 2015; Uegaki, 2021; Xu & Regier, 2014; Xu et al., 2016; Zaslavsky et al., 2018; Zaslavsky et al., 2021).

Here, we look towards a dynamic account of how meaning systems *come* to optimally balance simplicity and informativeness. We ask: what are some minimal conditions under which we can expect efficient meaning systems to emerge? Evolutionary game theory offers a framework for sketching an answer to this question. This paper investigates whether efficient compression is the typical result of evolution in a simple and well-studied model of meaning emergence.

Signaling Games

We model communication using signaling games (Lewis, 1969; Skyrms, 2010). In the games we consider, agents coordinate on signals to communicate about objects in a richly structured meaning space wherein atomic states bear similarity relations to one other via perceptual distance. A single round of the game models one communicative interaction as follows. Nature selects a state of the world to present to Sender. Sender observes this state and chooses a signal to send to Receiver, who cannot observe the state. Receiver then chooses an action to perform. Nature then awards payoff to both players if the action was appropriate. We can describe this game formally. Let $\mathcal{X}, \mathcal{Y}$ be two finite sets, $\Delta(\mathcal{X})$ denote the set of probability distributions over $\mathcal{X}$ and $\Delta(\mathcal{X})^{\mathcal{Y}}$ denote the set of conditional distributions of random variable $X \in \mathcal{X}$ given random variable $Y \in \mathcal{Y}$.

Definition. A *signaling game* is a tuple $\langle \mathcal{S}, \mathcal{W}, \hat{\mathcal{S}}, \sigma, \rho, u, P \rangle$ of states $\mathcal{S}$, signals $\mathcal{W}$, actions $\hat{\mathcal{S}}$, a Sender $\sigma(w|s) \in \Delta(\mathcal{W})^{\mathcal{S}}$, a Receiver $\rho(\hat{m}|w) \in \Delta(\hat{\mathcal{S}})^{\mathcal{W}}$, a utility function $u : \mathcal{S} \times \hat{\mathcal{S}} \to \mathbb{R}$, and a distribution over states $P(s) \in \Delta(\mathcal{S})$. Their joint payoff is given by:

$$\pi(\sigma, \rho) = \sum_s P(s) \sum_w \sigma(w|s) \sum_{\hat{s}} \rho(\hat{s}|w) \cdot u(s, \hat{s}) \quad (1)$$

$$= \mathbb{E}[u(S, \hat{S})], \quad (2)$$

To model a lexicon, we let $\hat{\mathcal{S}} = \mathcal{S}$ and talk only of states, which represent meanings. The payoff for Sender and Receiver can be thought of as an expected utility of a language for communication; rational communicative behavior is achieved if their signals maximize $\mathbb{E}[u(S, \hat{S})]$. Canonically, both players receive a payoff only if the action corresponds to the state. Payoff in this game rewards only perfect informativeness.

A similarity-maximizing game

The *similarity-maximizing* (sim-max) game weakens this latter assumption by endowing the meaning space with structure such that some meanings are closer 'guesses' than others (Jäger, 2007). That is, if the utility function u is classically defined to be the indicator function, in the sim-max game we let u depend on the perceptual similarity of states, $u \propto \text{sim}(s, \hat{s})$. Perceptual similarity calls for several modeling considerations. First, it should be defined over a set of objects standing in distance relations to each other. Second, less distance between states should correspond to higher similarity between them. Lastly, how quickly similarity decreases should depend on ***discriminative need*** —a pragmatic or environmental need to finely discriminate objects. Our particular formalizations are driven largely by a desire to make the minimal adjustments necessary to the simple Lewis-Skyrms model; scaling up to realistic settings, one would let the meaning space, models of agent perception, behavior, etc., be validated experimentally.

Meaning space We define the state space to be a set of integers, to model a situation in which states of the environment stand in minimally interesting physical distance relations.

$$\mathcal{S} = \hat{\mathcal{S}} = \{1, 2, \dots, n\} \quad (3)$$

Similarity function For the model of perceptual similarity, we follow Franke and Correia (2018) in using an independently motivated model from mathematical psychology (Nosofsky, 1986; Shepard, 1957). It is given by

$$\text{sim}_\alpha(x, y) = \exp\left(\frac{-(x-y)^2}{\alpha^2}\right), \quad (4)$$

where $\text{sim}_0(x, y)$ is the indicator function if $\alpha = 0$. Here α is an imprecision parameter. Total imprecision is reached as $\alpha \to \infty$, i.e. when all states are always rewarded; perfect discrimination between states is enforced when $\alpha = 0$, i.e. a classic Lewis-Skyrms game obtains. Note that the distance measure is squared difference between integers.

Semantic systems We measure the semantic system as a conditional distribution over Receiver interpreted meanings given Sender intended meanings:

$$p(\hat{s}|s) = \sum_w \sigma(w|s)\rho(\hat{s}|w), \quad (5)$$

where $p(\hat{s}|s)$ is the expected value of Receiver choosing meaning $\hat{s}$ given that Sender intended s, using their respective strategies for sending and receiving the available signals $\mathcal{W}$. Different distributions parameterized by Sender and Receiver behavior determine better or worse communicative success rates; we refer to these distributions interchangeably as communication/semantic systems and (generously) languages.

Rate-Distortion Theory

To evaluate whether languages in a sim-max game optimize the simplicity/informativeness trade-off, we apply the modeling framework introduced by Zaslavsky et al. (2018) to measure efficiency in terms of Rate-Distortion Theory (RDT), the branch of information theory concerned with optimizing lossy data compression (Berger, 1971; Cover & Thomas, 2006;

Shannon, 1959). In information-theoretic terms, languages minimize both *rate* (the resources necessary to compress a thought into a word, quantified in bits) and *distortion* (the error a listener makes in reconstructing a speaker's intention). We recapitulate a setting in which Sender meanings S are compressed into Receiver meanings $\hat{S}$, referring readers to Zaslavsky et al. (2018, SI, section 1.3) and Zaslavsky (2020) for the original motivation and formulation of this view of RDT and its formal connection to the Information Bottleneck framework (Harremoes & Tishby, 2007) as applied to studying semantic systems.[2]

Complexity The simplicity of a language is defined terms of its inverse, complexity. Formally, the complexity of the semantic system is given by the information rate:

$$I(S;\hat{S}) = \sum_{s} p(s) \sum_{\hat{s}} p(\hat{s}|s) \log \frac{p(\hat{s}|s)}{p(\hat{s})}, \tag{6}$$

where $p(s)$ is a prior over Sender states[3], $p(\hat{s}|s)$ is computed using Equation 5. The mutual information $I(S;\hat{S})$ quantifies the number of bits required, on average, to represent states of the environment S as $\hat{S}$. Minimal complexity at $I(S;\hat{S}) = 0$ is achieved when the players use one signal to represent every state of the environment. It can also be achieved when the players choose, for every state, each signal with equal probability. These latter strategies do not give anything useful to Receiver. To support successful communication, a language must allow some complexity, i.e. $I(S;\hat{S}) > 0$. Additional bits of complexity, however, require greater computational resources for speakers and listeners.

Communicative cost Informativeness is also defined in terms of its inverse, communicative cost, which we quantify by the expected distortion between Sender and Receiver:

$$\mathbb{E}[d(S,\hat{S})] = \sum_{s} p(s) \sum_{\hat{s}} p(\hat{s}|s) \cdot d(s,\hat{s}). \tag{7}$$

RDT can consider an arbitrary measure of distortion. Given that similarity of states in the game depends on their squared distance, we use $d(s,\hat{s}) = (s-\hat{s})^2$.

Bounds on efficiency In between the extremes of zero complexity and perfect accuracy, RDT predicts a continuum of functions $p(\hat{s}|s)$ that achieve the minimum distortion for a given complexity bound. These determine communicative strategies that optimally balance simplicity and informativeness. These functions are found by minimizing the Lagrangian

$$\mathcal{F}_{\beta}[p(\hat{s}|s)] = I(S;\hat{S}) + \beta \mathbb{E}[d(S,\hat{S})], \tag{8}$$

where β specifies the trade-off between complexity and communicative cost.[4] Equation 8 can be solved numerically by the Blahut-Arimoto algorithm (Arimoto, 1972; Blahut, 1972). It is an alternating minimization procedure that iterates the following equations until convergence:

$$p(\hat{s}) = \sum_{s} p(s) p(\hat{s}|s), \tag{9}$$

$$p(\hat{s}|s) = \frac{p(\hat{s}) \exp(-\beta \cdot d(s,\hat{s}))}{\sum_{\hat{s}} p(\hat{s}) \exp(-\beta \cdot d(s,\hat{s}))}. \tag{10}$$

The set of solutions to this problem define the Rate-Distortion curve, which in our setting represents the set of languages that perfectly optimize the simplicity/informativeness trade-off.

Predictions

If evolution leads to efficient semantic systems, then we should expect players' languages to converge to points close to the Rate Distortion curve. To model evolution, we perform two computational experiments: one involving Roth-Erev reinforcement learning, and another involving a discrete-time version of the standard replicator dynamic. There are a few reasons to use these dynamics, which we discuss in turn.

First, researchers have discovered that reinforcement learning and replicator dynamics operate across a variety of biological contexts (Glimcher, 2011; Schuster & Sigmund, 1983). Second, there are conditions under which both dynamics converge to Nash equilibria of Lewis-Skyrms games (Beggs, 2005; S. Huttegger et al., 2014); we intend our simulations to give intutions about potential relationships between equilibrium concepts in sim-max games and information theory. Analytical results about the connection between these two dynamics are limited; so as a robustness check in the present investigation, we apply both. Third, their simplicity lends generality. The dynamics we consider require relatively little cognitive sophistication, suggesting our model may capture features of efficient signaling in nature more generally. Lastly, although Roth-Erev learning and replicator dynamics have been deployed extensively to model the emergence of perfect coordination (Nowak & Krakauer, 1999; Skyrms, 2010; Spike et al., 2017) and vague meanings (Franke & Correia, 2018; O'Connor, 2014), to our knowledge they have not been proposed as models of the evolution of information theoretically efficient meaning systems.

Experiments

Roth-Erev reinforcement learning

One model for the evolution of meaning conventions in signaling games deploys a highly rudimentary form of reinforcement learning described by Erev and Roth (1998). This model is

[2]Future work will explore the robustness of our results measuring the trade-off in terms of the Information Bottleneck framework. This setting allows one to model perceptual uncertainty for Sender and Receiver, in addition to explicitly including the bottleneck variable $\mathcal{W}$ (signals) in the theoretical bound on efficiency. Meanwhile, see Tucker et al. (2022b, Appendix A) for the relation in RDT between KL-divergence and squared error.

[3]We report results for a uniform prior (communicative need distribution). We expect our results to be robust to some, but not all, possible need distributions. For discussion, see Komarova et al. (2007) in a related setting, and Barrett (2006), S. M. Huttegger (2007), and Skyrms (2010) in signaling games more generally.

[4]For precise definitions of Rate-Distortion functions and their optimization, see (Berger, 1971; Cover & Thomas, 2006).

popular in part because it requires minimal assumptions about the rationality of agents, and because it has been tractable enough for proving analytic results regarding convergence (Argiento et al., 2009; Skyrms, 2010). [5]

The Roth-Erev model of reinforcement learning specifies an *update* rule, which determines how agents' propensities towards pure strategies evolve over repeated plays of a game, and a *choice* rule, which determines the strategy that is played in a particular context. Let $q_{x,y}(t)$ denote an agent's propensity to choose strategy x in context y at time t, and r be the agent's reinforcement reward. The agent's propensity to play strategy x at the next time step is updated:

$$q_{x,y}(t+1) = q_{x,y}(t) + r \tag{11}$$

The choice rule (Luce, 1959) specifies the probability that an agent plays strategy x in context y, which is given by:

$$p_{x,y}(t) = \frac{q_{x,y}(t)}{\sum q_{x,y}(t)} \tag{12}$$

This model thus assumes that propensity weights and reinforcements are always positive. The reinforcement for Sender and Receiver when s is interpreted as $\hat{s}$ is defined:

$$r(s, \hat{s}) = u(s, \hat{s}) \cdot \eta, \tag{13}$$

where $\eta > 0$ is a positive constant acting as a learning rate.[6] This means that if w was the signal used during round t, the update to Sender's propensity $q_{s,w}(t)$ equals the update to Receiver's propensity $q_{w,\hat{s}}(t)$, which depends only on the similarity of s to $\hat{s}$. Only the strategies that were actually chosen by players are reinforced at each time step. A rudimentary form of reinforcement learning is modeled in the sense that successful actions become more likely over time.

To ensure that perfect accuracy is possible we report results for games with equal number of states and signals $|S| = |W| = n$. Since perfect communicative success becomes rare in large games, we set $n = 10$ (Barrett, 2006; S. M. Huttegger, 2007; Pawlowitsch, 2008).[7] We simulate multiple runs of Roth-Erev reinforcement learning for 10^6 rounds.

Discrete-time replicator dynamic

The standard replicator dynamic describes change in mean behavior in a population of game players. We imagine two distinct and virtually infinite populations of Senders and Receivers and model their evolution; the players are modeled as instantiating pure strategies, and behavioral strategies capture average population behavior. We follow Jäger (2007) and Franke and Correia (2018) in adopting a discrete-time version, which imagines that update steps are infinitesimally small. Formally, the updates to the behavioral strategies of a randomly sampled sender σ and a randomly sampled receiver ρ are given by:[8]

$$\sigma'(w|s) = \sigma(w|s) \cdot \sum_{\hat{s}} \rho(\hat{s}|w) \cdot u(s, \hat{s}) \tag{14}$$

$$\rho'(\hat{s}|w) = \rho(\hat{s}|w) \cdot \sum_{s} P(s) \cdot \sigma(s|w) \cdot u(s, \hat{s}) \tag{15}$$

In Equation 14, $\sigma(w|s)$ is the probability the average sender chooses w to communicate s. This represents the frequency of the corresponding type of pure strategy in the Sender population. Likewise, $\rho(\hat{s}|w)$ is the frequency of a strategy in the Receiver population. The frequency of an strategy evolves according to its current frequency and its fitness (i.e., the expected utility relative to the other population). Intuitively, this describes a situation of idealized Darwinian natural selection wherein 'like begets like'. Imitation and reproduction are two possible mechanisms that could implement this dynamic in theory. As in reinforcement learning, evolution is driven by the fact that successful strategies become more frequent. For comparability, we consider games with $|S| = |W| = 10$. We simulate the replicator dynamic for up to 200 time steps.

Results

We measure the efficiency of the emergent languages resulting from the two evolutionary dynamics in terms of distance to the Rate-Distortion curve. We report results for six different values of imprecision: $\alpha \in \{0, 1, 2, 4, 8, 16\}$ across 100 runs, varying random initialization of agents' strategies. In addition, to test whether convergence to a Rate-Distortion optimal system is trivial, i.e. whether an uninteresting random process could yield the diverse range of optimal solutions to the trade-off problem, we compare the evolved languages to many mathematically possible 'hypothetical' languages. We follow Zaslavsky et al. (2021) in generating hypothetical variants by randomly permuting the signals that a semantic system assigns to states. For each emergent system, we generate 100 such variants, resulting in a total 6000 hypothetical systems for each experiment.

Figure 1 displays the main results. Visual inspection suggests that both simple reinforcement learning agents and populations evolving according to the replicator dynamic do indeed approach the Rate-Distortion curve. Moreover, discriminative need in a game constrains the location of its emergent systems. The plots of evolutionary and learning trajectories in Figures 1C and 1D show that meaning systems in the sim-max game do traverse sub-optimal regions of the space before con-

[5] For discussion of the relationships between the many proposed learning mechanisms underlying emergent signaling, see Spike et al. (2017).

[6] The current results are for $\eta = 0.05$. The general trends presented here are relatively robust for small learning rates, e.g. those within $[0.05, 1.0]$. Note that the similarity function already has range $[0, 1]$. In the replicator dynamic, this scaling has no effect because frequencies are renormalized at each update.

[7] This limits the applicability of Roth-Erev learning to explanations of perfectly informative, large lexicons. Future work is needed to determine whether either the dynamics considered here can predict the emergence of Rate-Distortion optimal lexicons for larger meaning spaces.

[8] We refer readers to Franke and Correia (2018) for an explicit derivation of this formulation from the standard replicator equation.

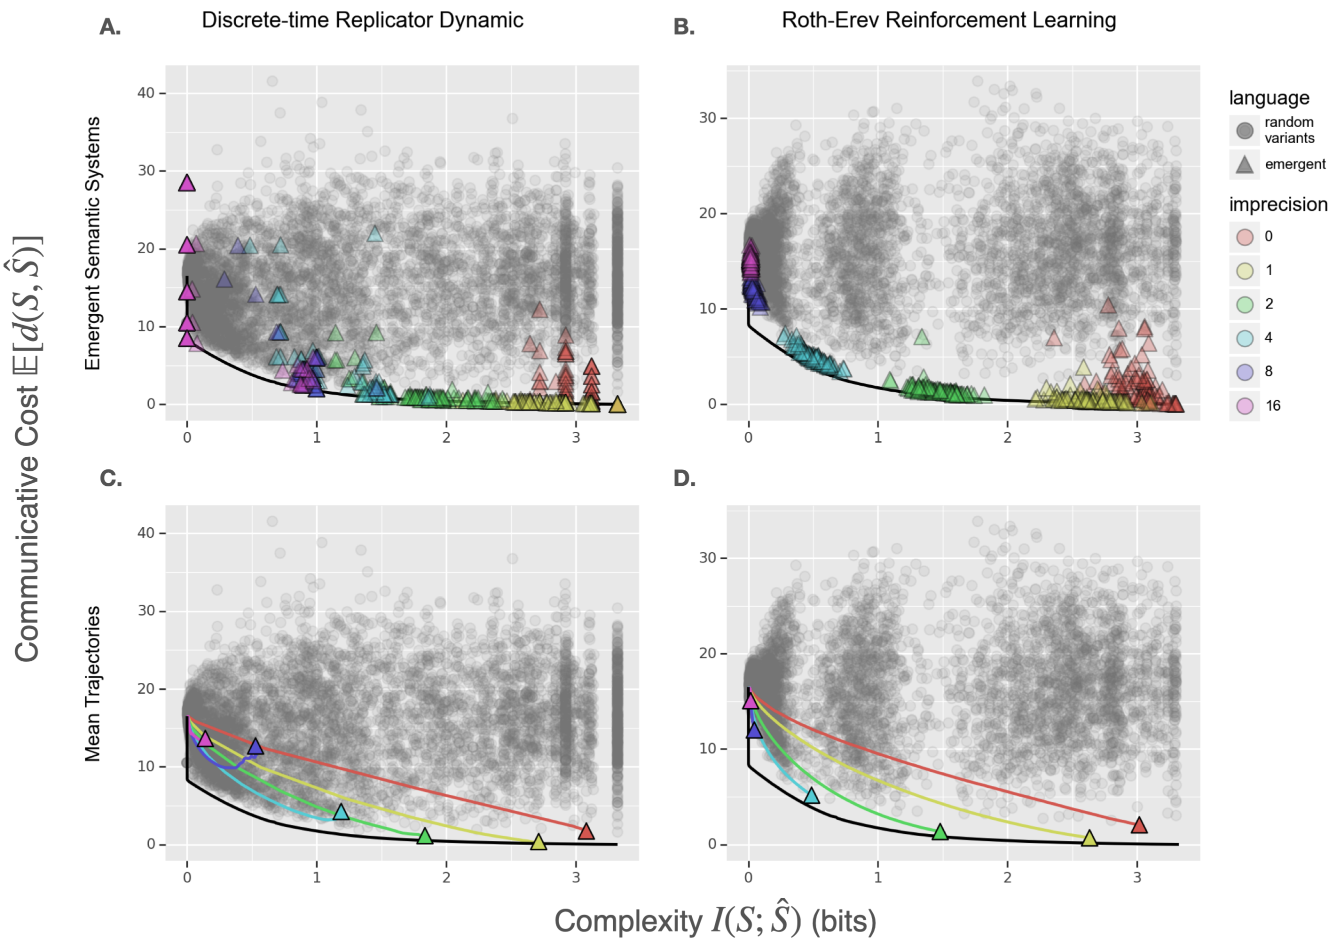

Figure 1: The evolution of information-theoretically efficient semantic systems in sim-max games with ten states and signals. The x-axis is complexity, given by the information rate $I(S; \hat{S})$ in bits. The y-axis is communicative cost, i.e. the expected squared difference between Sender and Receiver meanings (Equation 7). Black line: the Rate-Distortion curve, which is the set of optimal solutions to the Rate-Distortion objective of efficient compression. Triangles: the emergent semantic systems. Circles (gray): random variants of the emergent semantic systems, generated for comparison. Color: imprecision parameter α for the sim-max game that controls the utility of discriminating states. **A:** Systems (100 trials) evolved under the replicator dynamic. **B:** Systems (100 trials) learned by Roth-Erev agents. **C:** The mean evolutionary trajectories of systems during the replicator dynamic. **D:** Mean trajectories of systems during learning.

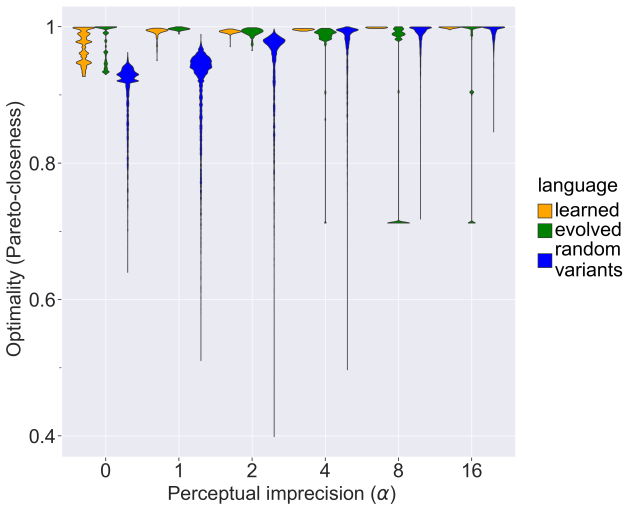

Figure 2: Distribution of semantic systems' efficiency across discriminative need. Orange: Roth-Erev learned systems. Green: systems evolved under the replicator dynamic. Blue: random variants of the emergent systems.

verging. Under this RDT-based analysis of sim-max games, inefficiency is achievable but often unstable in evolution.

However, efficiency is not always an expected result of either dynamics. Higher complexity systems are often far from the curve, and the replicator dynamic can yield very inefficient low-complexity systems. In the latter, large variation in resulting systems, depending on initial conditions, skews the mean trajectories. The high-complexity, inefficient systems may contain some of the well known 'partial pooling' equilibria of classic Lewis-Skyrms.[9] On the other hand, moderately high discriminative need typically results in near-optimality for both dynamics. Emergent systems also display a striking tendency to minimize complexity. Notably, Chaabouni et al. (2021) report that quantization from continous meanings to discrete symbols is sufficient for complexity minimization. This explanation alone is not sufficient our setting, since both our meanings and signals are discrete; we leave a full explanation of this effect to future work.

A quantitative comparison between the efficiency of languages emerging from Roth-Erev learning and the replicator dynamic vs. the set of pooled hypothetical variants is depicted in Figure 2. The efficiency of a language is measured by the normalized minimum Euclidean distance to its closest point on the Rate-Distortion curve. This plot shows that pressures for utility/fitness in a simple model of communication are sufficient for the evolution information-theoretically efficient meaning systems, and that (with the exception of high imprecision settings) merely random processes do not usually deliver efficiency.

[9]There are analytic results about the convergence to these non-strict Nash equilibria under Roth-Erev learning (Hu et al., 2011) and replicator dynamics (S. M. Huttegger, 2007; Pawlowitsch, 2008).

Summary

In simulations, signaling agents in sim-max games converge to communicative behavior that is near-optimal from the standpoint of Rate-Distortion Theory. Additionally, the degree to which perceptual distinctions are rewarded in the game typically constrains where, along the RD curve, languages converge. This suggests that when there is joint environmental pressure to discriminate stimuli and to coordinate on signals, even rudimentary adaptive dynamics can carry semantic systems to information-theoretic efficiency.

Related work

We aim to contribute to the existing literature characterizing the evolution of communicative efficiency. This includes dynamic computational and evolutionary models of perceptually-based color categorization (Komarova et al., 2007; Steels & Belpaeme, 2005; Zuidema & Westermann, 2003), iterated learning models invoked to explain the emergence of simple and informative lexicons (Carr et al., 2020; Carstensen et al., 2014; Kirby et al., 2015), and the emergence of human-like semantic systems achieving efficient compression from reference games in deep reinforcement learning settings (Carlsson et al., 2021; Chaabouni et al., 2021; Tucker et al., 2022a, 2022b). In addition, our current work is continuous with research that explores how convex meanings and vague terms can evolve in sim-max games (Correia & Franke, 2019; Franke & Correia, 2018; Jäger, 2007; Jäger & van Rooij, 2007; O'Connor, 2014).

Our model is distinct in connecting (i) views in which natural language lexicons result from optimally balancing cognitive complexity and communicative accuracy to (ii) evolutionary game theoretic approaches describing how meaning systems can emerge naturally from minimally sophisticated agents. To our knowledge, this paper is the first to unify these two adjacent literatures.

Importantly, we do not intend to provide a dynamic framework that scales to complex, naturalistic settings, or an empirical account of human category learning or cultural evolution. Instead, we believe our goals align with those of other researchers in the evolution of language, especially Komarova et al. (2007), who intend to "demonstrate what can be achieved using only the most rudimentary forms of [category] observation and communication together with an elementary evolutionary dynamics." Additionally, we hope to provide intuitions for potential mathematical connections between evolutionary game theory and information theory. We look forward to extending the current results in these directions in future work.

Acknowledgments

We thank Jeffrey Barrett, Pedro Correia, Michael Franke, Richard Futrell, Brian Skyrms, Cailin O'Connor, Shane Steinert-Threlkeld, Noga Zaslavsky and several anonymous reviewers for valuable feedback.

References

Argiento, R., Pemantle, R., Skyrms, B., & Volkov, S. (2009). Learning to signal: Analysis of a micro-level reinforcement model. *Stochastic Processes and their Applications*, *119*(2), 373–390.

Arimoto, S. (1972). An algorithm for computing the capacity of arbitrary discrete memoryless channels [Conference Name: IEEE Transactions on Information Theory]. *IEEE Transactions on Information Theory*, *18*(1), 14–20.

Barrett, J. A. (2006). *Numerical Simulations of the Lewis Signaling Game: Learning Strategies, Pooling Equilibria, and the Evolution of Grammar* (tech. rep.). Institute for Mathematical, Behavioral Sciences.

Beggs, A. W. (2005). On the convergence of reinforcement learning. *Journal of Economic Theory*, *122*(1), 1–36.

Berger, T. (1971). *Rate distortion theory; a mathematical basis for data compression.* Prentice-Hall
Includes bibliographical references (pages 293-301).

Blahut, R. (1972). Computation of channel capacity and rate-distortion functions [Conference Name: IEEE Transactions on Information Theory]. *IEEE Transactions on Information Theory*, *18*(4), 460–473.

Carlsson, E., Dubhashi, D., & Johansson, F. D. (2021). Learning Approximate and Exact Numeral Systems via Reinforcement Learning. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *43*(43).

Carr, J. W., Smith, K., Culbertson, J., & Kirby, S. (2020). Simplicity and informativeness in semantic category systems. *Cognition*, *202*, 104289.

Carstensen, A., Xu, J., Smith, C. T., & Regier, T. (2014). Language evolution in the lab tends toward informative communication. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *36*(36).

Chaabouni, R., Kharitonov, E., Dupoux, E., & Baroni, M. (2021). Communicating artificial neural networks develop efficient color-naming systems [Publisher: Proceedings of the National Academy of Sciences]. *Proceedings of the National Academy of Sciences*, *118*(12), e2016569118.

Correia, J. P., & Franke, M. (2019). Towards an Ecology of Vagueness. In R. Dietz (Ed.), *Vagueness and Rationality in Language Use and Cognition* (pp. 87–113). Springer International Publishing.

Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience.

Denić, M., Steinert-Threlkeld, S., & Szymanik, J. (2022). Indefinite Pronouns Optimize the Simplicity/Informativeness Trade-Off. *Cognitive Science*, *46*(5), e13142
_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cogs.13142.

Erev, I., & Roth, A. E. (1998). Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria [Publisher: American Economic Association]. *The American Economic Review*, *88*(4), 848–881.

Franke, M., & Correia, J. P. (2018). Vagueness and Imprecise Imitation in Signalling Games [Publisher: The University of Chicago Press]. *The British Journal for the Philosophy of Science*, *69*(4), 1037–1067.

Glimcher, P. W. (2011). Understanding dopamine and reinforcement learning: The dopamine reward prediction error hypothesis [Publisher: Proceedings of the National Academy of Sciences]. *Proceedings of the National Academy of Sciences*, *108*(supplement_3), 15647–15654.

Harremoes, P., & Tishby, N. (2007). The Information Bottleneck Revisited or How to Choose a Good Distortion Measure [ISSN: 2157-8117]. *2007 IEEE International Symposium on Information Theory*, 566–570.

Hofbauer, J., & Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press.

Hu, Y., Skyrms, B., & Tarrès, P. (2011). Reinforcement learning in signaling game [arXiv:1103.5818 [math]]
Comment: 6 figures.

Huttegger, S., Skyrms, B., Tarrès, P., & Wagner, E. (2014). Some dynamics of signaling games [Publisher: Proceedings of the National Academy of Sciences]. *Proceedings of the National Academy of Sciences*, *111*(supplement_3), 10873–10880.

Huttegger, S. M. (2007). Evolution and the Explanation of Meaning [Publisher: Cambridge University Press]. *Philosophy of Science*, *74*(1), 1–27.

Imel, N., & Steinert-Threlkeld, S. (2022). Modal semantic universals optimize the simplicity/informativeness tradeoff. *Proceedings of semantics and linguistic theory (SALT 32)*.

Jäger, G. (2007). The evolution of convex categories. *Linguistics and Philosophy*, *30*(5), 551–564.

Jäger, G., & van Rooij, R. (2007). Language structure: Psychological and social constraints. *Synthese*, *159*(1), 99–130.

Kågebäck, M., Carlsson, E., Dubhashi, D., & Sayeed, A. (2020). A reinforcement-learning approach to efficient communication [Publisher: Public Library of Science]. *PLOS ONE*, *15*(7), e0234894.

Kemp, C., & Regier, T. (2012). Kinship categories across languages reflect general communicative principles. *Science (New York, N.Y.)*, *336*(6084), 1049–1054.

Kemp, C., Xu, Y., & Regier, T. (2018). Semantic typology and efficient communication. *Annual Review of Linguistics*, 1–23.

Kirby, S., Tamariz, M., Cornish, H., & Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition*, *141*, 87–102.

Komarova, N. L., Jameson, K. A., & Narens, L. (2007). Evolutionary models of color categorization based on discrimination. *Journal of Mathematical Psychology*, *51*(6), 359–382.

Levinson, S. C. (2012). Kinship and Human Thought [_eprint: https://www.science.org/doi/pdf/10.1126/science.1222691]. *Science*, *336*(6084), 988–989.

Lewis, D. K. (1969). *Convention: A philosophical study*. Cambridge, MA, USA: Wiley-Blackwell.

Luce, R. D. (D. (1959). *Individual choice behavior; a theoretical analysis.* Wiley
Includes bibliographical references (pages 147-149).

Mollica, F., Bacon, G., Zaslavsky, N., Xu, Y., Regier, T., & Kemp, C. (2021). The forms and meanings of grammatical markers support efficient communication [Publisher: Proceedings of the National Academy of Sciences]. *Proceedings of the National Academy of Sciences*, *118*(49), e2025993118.

Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship [Place: US Publisher: American Psychological Association]. *Journal of Experimental Psychology: General*, *115*, 39–57.

Nowak, M. A., & Krakauer, D. C. (1999). The evolution of language [Publisher: Proceedings of the National Academy of Sciences]. *Proceedings of the National Academy of Sciences*, *96*(14), 8028–8033.

O'Connor, C. (2014). The Evolution of Vagueness. *Erkenntnis*, *79*(4), 707–727.

Pawlowitsch, C. (2008). Why evolution does not always lead to an optimal signaling system. *Games and Economic Behavior*, *63*(1), 203–226.

Regier, T., Kemp, C., & Kay, P. (2015). Word meanings across languages support efficient communication Informativeness and simplicity as competing principles. In B. MacWhinney & W. O'Grady (Eds.), *The Handbook of Language Emergence* (pp. 237–263). Wiley.

Schuster, P., & Sigmund, K. (1983). Replicator dynamics. *Journal of Theoretical Biology*, *100*(3), 533–538.

Shannon, C. E. (1959). Coding Theorems for a Discrete Source With a Fidelity Criterion. *Claude E. Shannon: Collected Papers* (pp. 325–350). Institute of Radio Engineers, International Convention Record.

Shepard, R. N. (1957). Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space. *Psychometrika*, *22*(4), 325–345.

Skyrms, B. (2010). *Signals: Evolution, Learning, and Information*. Oxford University Press.

Spike, M., Stadler, K., Kirby, S., & Smith, K. (2017). Minimal Requirements for the Emergence of Learned Signaling [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cogs.12351]. *Cognitive Science*, *41*(3), 623–658.

Steels, L., & Belpaeme, T. (2005). Coordinating perceptually grounded categories through language: A case study for colour [Publisher: Cambridge University Press]. *Behavioral and Brain Sciences*, *28*(4), 469–489.

Tucker, M., Levy, R. P., Shah, J., & Zaslavsky, N. (2022a). Generalization and Translatability in Emergent Communication via Informational Constraints. *NeurIPS 2022 Workshop on Information-Theoretic Principles in Cognitive Systems*.

Tucker, M., Levy, R. P., Shah, J., & Zaslavsky, N. (2022b). Trading off utility, informativeness, and complexity in emergent communication. In A. H. Oh, A. Agarwal, D. Belgrave, & K. Cho (Eds.), *Advances in neural information processing systems*.

Uegaki, W. (2021). The informativeness / complexity trade-off in the domain of Boolean connectives. *Linguistic Inquiry*.

Xu, Y., & Regier, T. (2014). Numeral systems across languages support efficient communication: From approximate numerosity to recursion. *Cognitive science society (cogsci-2014)*, 1802–1807.

Xu, Y., Regier, T., & Malt, B. C. (2016). Historical Semantic Chaining and Efficient Communication: The Case of Container Names. *Cognitive Science*, *40*(8), 2081–2094.

Zaslavsky, N. (2020). *Information-Theoretic Principles in the Evolution of Semantic Systems* (Ph.D. Thesis). The Hebrew University of Jerusalem.

Zaslavsky, N., Garvin, K., Kemp, C., Tishby, N., & Regier, T. (2019). Evolution and efficiency in color naming: The case of Nafaanra. *CogSci*, 68.

Zaslavsky, N., Kemp, C., Regier, T., & Tishby, N. (2018). Efficient compression in color naming and its evolution. *Proceedings of the National Academy of Sciences*, *115*(31), 7937–7942.

Zaslavsky, N., Maldonado, M., & Culbertson, J. (2021). Let's talk (efficiently) about us: Person systems achieve near-Optimal compression. *Proceedings of the 43rd Annual Meeting of the Cognitive Science Society*.

Zuidema, W., & Westermann, G. (2003). Evolution of an Optimal Lexicon under Constraints from Embodiment. *Artificial Life*, *9*(4), 387–402.