

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Convention-formation in iterated reference games

Permalink

<https://escholarship.org/uc/item/46r654df>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 39(0)

Authors

Hawkins, Robert X.D

Frank, Michael C.

Goodman, Noah D.

Publication Date

2017

Peer reviewed

Convention-formation in iterated reference games

Robert X.D. Hawkins, Michael C. Frank, Noah D. Goodman

{rxdh, mcfrank, ngoodman}@stanford.edu

Department of Psychology, Stanford University

Abstract

What cognitive mechanisms support the emergence of linguistic conventions from repeated interaction? We present results from a large-scale, multi-player replication of the classic *tangrams task*, focusing on three foundational properties of conventions: arbitrariness, stability, and reduction of utterance length over time. These results motivate a theory of convention-formation where agents, though initially uncertain about word meanings in context, assume others are using language with such knowledge. Thus, agents may learn about meanings by reasoning about a knowledgeable, informative partner; if all agents engage in such a process, they successfully coordinate their beliefs, giving rise to a conventional communication system. We formalize this theory in a computational model of language understanding as social inference and demonstrate that it produces all three properties in a simplified domain.

Keywords: conventions; pragmatics; communication

Introduction

Just as drivers depend on shared behavioral conventions to safely navigate traffic, successful communication depends on a set of shared linguistic conventions. Speakers of different languages around the world refer to the same object in many different ways, yet when ordering a coffee in San Francisco, one can confidently use the English word “coffee” and be understood. How do these conventions – classically characterized by Lewis (1969) as arbitrary but stable solutions to recurring coordination problems – form in the first place?

While *global* conventions adopted and sustained throughout a large population of speakers may develop over longer time scales, we also effortlessly coordinate on *local* conventions – or conceptual pacts (Brennan & Clark, 1996) – within the span of a single dialogue. For example, when discussing possible conditions to use in an upcoming experiment, a team of collaborators might begin the meeting using long descriptions to refer to each condition but end the meeting using conventional terms like “condition A” and “condition B.” Since global conventions are hypothesized to emerge through diffuse, repeated interactions of this more local kind (Garrod & Doherty, 1994), the cognitive mechanisms underlying convention-formation in such interactions are of foundational interest.

In a seminal study by Krauss & Weinheimer (1964), pairs of participants played a cooperative language game where they were presented with arrays of ambiguous shapes in randomized orders. The players were assigned the roles of *director* and *matcher* and allowed to talk freely. The matcher’s goal was to rearrange their shapes to match the director’s board, and the director’s goal was to communicate useful descriptions. Over multiple rounds, descriptions were dramatically shortened: an early description like “upside-down martini glass in a wire stand,” became simply “martini” by the

end. Later studies (e.g. Clark & Wilkes-Gibbs, 1986) refined this paradigm, using larger arrays of tangram-like figures and emphasizing the intricate back-and-forth process through which speakers and listeners negotiate over references. The referring expressions generated by participants in these studies revealed a number of rich qualitative phenomena. Here, we focus on three that are both prescribed top-down by theories of convention-formation and also arise bottom-up as major axes of variation in our data: arbitrariness, stability, and the systematic reduction of utterance length over time.

Arbitrariness is a definitional property of conventions (Lewis, 1969): there must be multiple solutions that would be equally successful as long as both players “agree” (e.g. driving on the left vs. right side of the road). By the final round in a language game, for example, one pair might successfully use the expression ‘dancer’ to refer to a tangram, while another might use ‘skater’. The other definitional property we consider is *stability*: it is in everyone’s best interest to keep using a convention once established. Finally, *reduction* is more specific to the reference game paradigm and refers to the transformation of longer, complex expressions into simpler expressions over the course of interaction, as Krauss & Weinheimer (1964) observed. While this broad phenomenon has been replicated many times, exactly what is reduced remains an open empirical question.

Theories of convention-formation differ primarily in the extent to which sophisticated social reasoning and common ground is required. At one extreme, agents use simple heuristic updating rules and do not need to represent or reason about other agents at all (Barr, 2004; Centola & Baronchelli, 2015; Young, 2015). Simulations elegantly show how arbitrary signaling systems can spread and come to stably dominate large populations. However, due to their ‘rich get richer’ dynamic, it is not clear how simple heuristic updating mechanisms alone could account for reduction in repeated interaction. At the other extreme are theories in which agents recursively track what information is *mutual knowledge*, often formalized in a game theoretic setting (Lewis, 1969). Wilkes-Gibbs & Clark (1992) and others have proposed that agents engage in a collaborative process of actively establishing mutual knowledge, though the mechanisms allowing conventions to emerge under such conditions have not been instantiated in a formal model to our knowledge.

In this paper, we argue for a theoretical position on the spectrum between these poles: *conventions form when uncertain agents treat their partners’ knowledge as ground truth*. In other words, agents assume their partner is knowledgeable and rationally using some conventional lexicon mapping labels to meanings but are themselves initially unsure of its

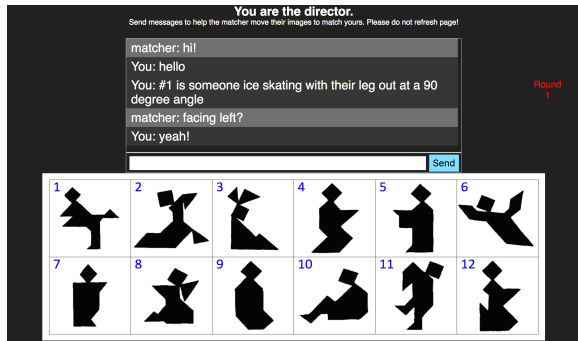


Figure 1: Example trial in experimental interface. Both players could freely use the chat box, and the matcher could click and drag the tangram images.

identity. Through observing their partner’s behavior in repeated actions, agents learn and adopt that lexicon, though their partner in fact begins in the same state of ignorance. When both agents independently adopt such a social learning strategy, they align to one another, coordinating on and implicitly creating shared conventions.

To motivate this theory, we first conduct a large-scale replication of the tangrams task on the web, which has traditionally been limited to relatively small sample sizes in the lab. We use distributions of lexical and syntactic features in the text corpus to operationalize arbitrariness, stability, and reduction, which have been difficult to analyze at a fine-grained level due to the sparseness of existing data. Taking these insights into account, we then formalize our theory in a computational model of communication in repeated reference games based on recent successes capturing language understanding as social inference (Goodman & Frank, 2016; Goodman & Stuhlmiller, 2013). Finally, we show that this model qualitatively produces all three empirical signatures in a simplified domain inspired by the tangrams task.

Replication of the Tangrams task

To collect a corpus of reference game dialogue that supports more detailed analyses of convention-formation, we ported the tangrams task used in Clark & Wilkes-Gibbs (1986) to a real-time, multi-player web environment.

Methods

Participants 200 participants were recruited from Amazon’s Mechanical Turk and paired into dyads to play a real-time communication game using the framework in Hawkins (2015). We excluded games that terminated before the completion of 6 rounds and where participants reported a native language different from English, leaving a corpus of 67 complete games with a total of 9967 utterances.

Stimuli On every trial of the game, both participants were shown a 6×2 grid containing twelve tangram shapes, reproduced from Clark & Wilkes-Gibbs (1986). Cells were labeled with fixed numbers from one to twelve in order to help par-

ticipants easily refer to locations in the grid (see Fig. 1).

Procedure After passing a short quiz about task instructions, participants were randomly assigned the role of either ‘director’ or ‘matcher’ and automatically paired into virtual rooms containing a chat box and grid of stimuli. Both participants could freely use the chat box to communicate at any time. The director’s tangrams were fixed in place, but the matcher could click and drag the shapes to reorder them. The director had to send messages about the locations of different tangrams on their fixed board; the matcher had to identify the corresponding tangram shapes and move them to the correct locations. When the players were satisfied that their boards matched, the matcher clicked a ‘submit’ button that gave players feedback on their score (out of 12) and scrambled the tangrams for the next round. After six rounds, players were redirected to a short exit survey. We collected the raw text of every message sent and every swapping action taken by the matcher¹.

Results

Arbitrariness and stability We begin by examining signatures of *arbitrariness* and *stability* in our data. We operationalize these concepts using the information-theoretic measure of entropy:

$$H(W) = \sum_w P(w) \log P(w)$$

where $P(w)$ denotes the distribution over frequencies of word tokens used within a game. Broadly speaking, entropy measures the predictability of a distribution. It is maximized when all elements are equally likely and declines as the distribution becomes more structured, i.e. when the probability mass is concentrated on a small subset of elements.

To derive predictions, we consider a permutation-test null model in which utterances on each of the six rounds are scrambled across games, designed to break any existing structure in each game’s idiosyncratic word distributions. The mean empirical entropy should only differ from the null distribution of entropies generated from this scrambling process if *both* arbitrariness and stability hold.

First, note that if stability did *not* hold, scrambling would have no effect on the entropy within individual games: a given speaker would already use different words each round, and swapping out the identity of those words would not affect the entropy of the word distribution. There would be no structure to break.

If stability holds but arbitrariness does not, all players would adopt the single optimal (non-arbitrary) way to refer to each tangram. Therefore, the entropy of their word distributions also should not be affected by scrambling: a speaker’s real words would be swapped out for the same tokens generated by another speaker. Scrambling wouldn’t break the

¹Data are available at <https://cocolab.stanford.edu/datasets/tangrams.html>

	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10
unigrams	a	like	looks	the	one	with	to	of	and	on
bigrams	looks like	like a	a person	is a	to the	with a	the right	the left	the one	a square

Table 1: Top 10 unigrams and bigrams with the highest reduction




Figure 2: Reduction phenomena. From left: (1) mean message length in words per tangram, (2) mean number of listener messages, (3) proportion of utterances containing adjectival clauses, (4) proportion of utterances containing subordinate clauses. Error bars are bootstrapped 95% CIs.

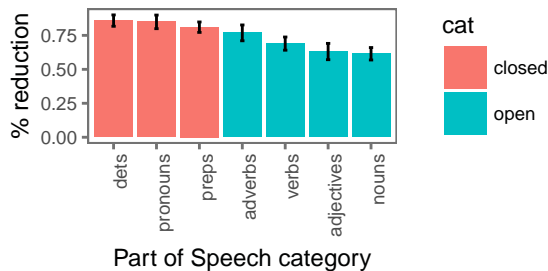


Figure 3: Reduction rates for different parts of speech. Error bars are bootstrapped 95% CIs.

structure of the distribution, because the structure would be the same for all participants.

Finally, if both arbitrariness and stability hold, then different speakers would adopt different referring expressions that persist from round to round. Hence, scrambling should *increase* the average game’s entropy from a relatively low level: each game’s idiosyncratic, concentrated distribution of words would be mixed together to form more heterogeneous and therefore high-entropy distributions.

To test this prediction, we computed the average within-game entropy for 1000 different permutations of speaker utterances. Since this permutation scheme keeps the number of messages per participant constant and simply swaps out the content of those messages within each round, it controls for the fact that some speakers sent more messages than others and also that speakers in earlier rounds use more words (see next section). We found that our null distribution lay within the interval [4.88, 4.91], which is significantly higher than the true entropy (averaged across games) of 4.36, $p < 0.001$. This pattern is consistent only with signatures of both arbitrariness and stability.

Reduction Next, we turn to a set of analyses examining reduction in utterance length over the course of the experiment. At the coarsest level, we find that the mean number of

words used by speakers decreases over time (see Fig. 2). This decrease replicates a highly reliable reduction effect found throughout the literature on iterated reference games (Brennan & Clark, 1996; Krauss & Weinheimer, 1964). Likely due to our purely textual (vs. spoken) interface, participants in our task used significantly fewer words overall than previously reported (e.g. an average of 20 words on the 1st round, compared to 40 in Clark & Wilkes-Gibbs (1986)) The following analyses break down this broad reduction into a finer-grained set of phenomena.

The next level of granularity motivating our model approach concerns which kinds of words are most likely to be dropped. Is the speaker adopting a shorthand where they drop uninformative function words, or are they simplifying or narrowing their descriptions by omitting meaningful details (Clark & Wilkes-Gibbs, 1986)? We used the Stanford CoreNLP part-of-speech tagger (Toutanova, Klein, Manning, & Singer, 2003) to count the number of words belonging to each part of speech in each message. Fig. 3 shows the percent reduction of different parts of speech from the first round to the sixth round. We find that determiners (‘the’, ‘a’, ‘an’) are the most likely class of words to be dropped with an 86% reduction rate, on average. Nouns (‘dancer’, ‘rabbit’) are the least likely class to be dropped with only an 62% rate. Closed-class parts of speech are strictly more likely to be dropped than open-class parts of speech.

While this finding is consistent with the possibility that speakers adopt a shorthand using more fragments as the game proceeds, we find a more complex dynamic by examining the table of unigrams and bigrams most likely to be dropped (see Table 1). Note that alongside dropped articles (‘a’, ‘the’), there are a number of words that form conjunctions (‘and’) and modifiers (‘of’, ‘with’, ‘the right’). In other words, it may be more likely that when function words are dropped, it is primarily as part of larger grammatical units that provide additional information in identifying the target.

We explicitly examined this hypothesis by running the Stanford constituency parser (Schuster & Manning, 2016),

tagging the occurrence of subordinate/adverbial clauses (‘sitting *facing left*’) and adjectival clauses (‘angel *that is praying*’).² We found that both were reduced over the course of the game (see Fig. 2), lending additional support for the hypothesis that whole meaningful clauses are increasingly omitted. This result prompts a characterization of reduction where, due to uncertainty at the outset about the usefulness of any particular lexical unit, initial phrases pile on multiple partially redundant modifiers and descriptors. As the game progresses and ambiguity of reference decreases, these additional meaningful units become less useful and can be dropped. We return to this characterization more formally within the scope of our model below.

Model

Here, we present a probabilistic model of language production under uncertainty, which captures several of the signature properties of conventions shown above. This model belongs to the family of Rational Speech Act (RSA) models, which have been successful in explaining a wide range of linguistic phenomena – including scalar implicature, adjectival vagueness, overinformativeness, indirect questions, and non-literal language use – as arising from a process of recursive social reasoning. Most previous applications of RSA have focused on the listener’s problem of language comprehension, but the puzzle of conventionalization is primarily a question of speaker production. An n th order pragmatic speaker trying to convey a particular state of affairs $s \in \mathcal{S}$ assuming lexicon \mathcal{L} is assumed to select an utterance $u \in \mathcal{U}$ by trading off its expected informativity (with respect to a rational listener agent) against its cost, usually based on length (Goodman & Frank, 2016):

$$S_n(u|s, \mathcal{L}) \propto \exp(\alpha \log L_{n-1}(s|u, \mathcal{L}) - \text{cost}(u))$$

where α is a soft-max optimality parameter controlling the extent to which the speaker maximizes over listener informativity. The listener, in turn, inverts the speaker model to reason about what underlying state s the speaker is trying to convey, given their utterance u :

$$L_n(s|u, \mathcal{L}) \propto P(s)S_n(u|s, \mathcal{L})$$

This recursion bottoms out in a *literal listener* who directly looks up the meaning of the utterance in the lexicon:

$$L_0(s|u, \mathcal{L}) \propto \mathcal{L}(u, s) \cdot P(s)$$

As in several other recent applications of RSA (Graf, Degen, Hawkins, & Goodman, 2016), we use a graded semantics, where utterances are better or worse descriptions of particular referents. For instance, the utterance “dancer” may initially be expected to apply to a photorealistic image of a ballerina ($\mathcal{L}(\text{‘dancer’}, \textit{ballerina}) = 0.99$) more than an abstract

²Specifically, we used the Universal Dependencies tags `csubj`, `ccomp`, `xcomp`, and `advcl` for subordinate clauses and `acl` for adjectival clauses (Schuster & Manning, 2016)

image of one ($\mathcal{L}(\text{‘dancer’}, \textit{abstract ballerina}) = 0.6$), but apply to both better than a non-category member like an image of a dog ($\mathcal{L}(\text{‘dancer’}, \textit{dog}) = 0.05$).

Our approach to convention-formation begins with the additional assumption of *lexical uncertainty* (Bergen, Levy, & Goodman, 2016; Smith, Goodman, & Frank, 2013). In other words, we assume that instead of having perfect knowledge of \mathcal{L} , the listener has uncertainty over the exact meanings of lexical items in the current context (i.e. it is initially unclear which of the ambiguous tangram shapes “the dancer” might refer to). They begin with some prior $P(\mathcal{L})$ about the identity their partner’s true lexicon, which may be initially biased toward certain meanings. By conditioning on repeated observations of their partner’s behavior, they use Bayes rule to infer this true lexicon:

$$P_{L_n}(\mathcal{L}|d) \propto P(\mathcal{L}) \prod_i S_n(s_i|u_i, \mathcal{L})$$

where $d = \{s_i, u_i\}$ is a set of observations of s_i and u_i coming from previous exchanges³. The listener marginalizes over this posterior when interpreting the speaker’s utterance:

$$L_n(s|u, d) \propto \sum_{\mathcal{L}} P_{L_n}(\mathcal{L}|d) L_n(s|u, \mathcal{L})$$

The speaker, in turn, considers what utterances would be most informative for such a listener:

$$S_n(u|s, d) \propto \exp(\alpha \log \left(\sum_{\mathcal{L}} P_{S_n}(\mathcal{L}|d) L_{n-1}(s|u, \mathcal{L}) \right) - \text{cost}(u))$$

where the posterior over lexica $P_{S_n}(\mathcal{L}|d)$, uses the listener likelihood L_{n-1} . For the purposes of this paper, we fix the depth of recursion at $n = 2$. This model is implemented in the probabilistic programming language WebPPL (Goodman & Stuhlmiller, electronic).⁴ Following Smith et al. (2013), we begin by showing how a random initial choice is taken to be evidence for a particular lexicon and becomes the base for successful communication even though neither party knows its meaning at the outset.

Results

Arbitrariness and stability Consider an environment with two abstract shapes ($\{s_1, s_2\}$), where the speaker must choose between two utterances ($\{u_1, u_2\}$) incurring equal cost. Their prior $P(\mathcal{L})$ over the meaning of each utterance is given by a Beta distribution⁵, so on the first round both utterances are

³There is a broader debate over the timescales at which lexicons and lexicon learning mechanisms operate; here, we assume a discourse-level structure to the lexicon, where there is uncertainty over how words are used in the given conversation. See Frank, Goodman, & Tenenbaum (2009) for a related approach at the scale of cross-situational word learning.

⁴All results can be reproduced running our code in the browser at <http://forestdb.org/models/conventions.html>

⁵In our implementation, we enumerate over coarse-grained bins; preliminary experiments using variational inference on the full continuous distribution give similar results.

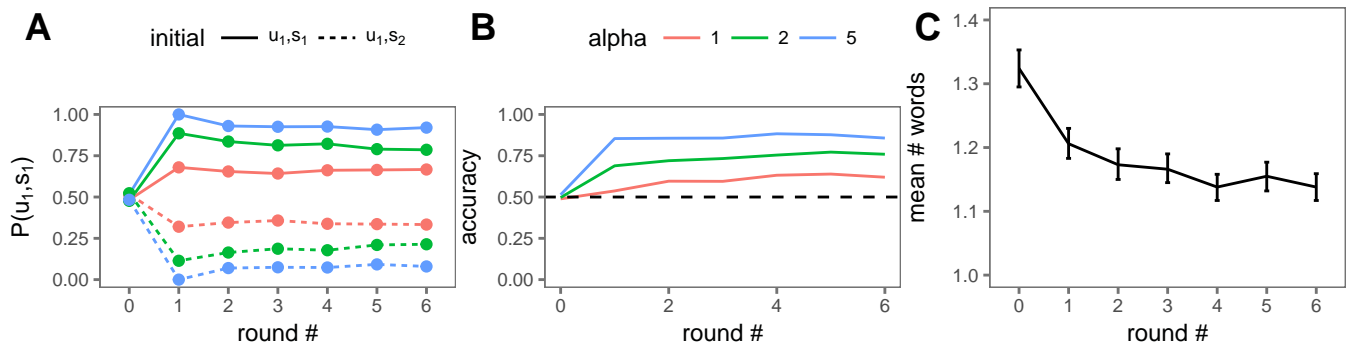


Figure 4: (A) Probability of speaker using u_1 to refer to s_1 , broken out by initial observation: while players are initially ambivalent between the two labels (arbitrariness), the initial mapping is likely to persist (stability). (B) Accuracy rises as speaker and listener align. (C) When conjunctions are introduced into the grammar, utterances get shorter over time (reduction).

equally likely to apply to either shape. If the speaker was trying to get their partner to pick s_1 , then, since each utterance is equally (un)informative, they would randomly sample one (say, u_1), and observe the listener’s selection of a shape (say, s_1). On the next round, the speaker uses the observed pair $\{u_1, s_1\}$ to update their beliefs about their partner’s true lexicon, uses these beliefs to generate a new utterance, and so on. To examine expected dynamics over multiple rounds, we forward sample many possible trajectories.

We observe several important qualitative effects in our simulations. First, the fact that a knowledgeable listener responds to utterance u with s provides evidence for lexicons in which u is a good fit for s , hence the likelihood of the speaker using u to refer to s increases on subsequent rounds (see Fig.4A). In other words, the initial symmetry between the meanings can be broken by initial random choices, leading to completely *arbitrary but stable mappings* in future rounds. Second, because the listener is also learning the lexicon from these observations under the same set of assumptions, they converge on a shared set of meanings; hence, expected *accuracy* rises on future rounds (see Fig. 4B). Third, because one’s partner is assumed to be pragmatic, agents can also learn about *unheard* utterances. Observing $d = \{u_1, s_1\}$ also provides evidence that u_2 is *not* a good fit for s_1 by Gricean maxims: if u_2 were a better fit for s_1 , the speaker would have used it instead (Grice, 1975). Finally, *failed references* lead to conventions just as effectively as successful references: if the speaker intends s_1 and says u_1 , but then the listener incorrectly picks s_2 , the speaker will take this as evidence that u_1 actually means s_2 in their partner’s lexicon and become increasingly likely to use it that way on subsequent rounds.

Reduction in utterance length Finally, we show how our model explains reduction of utterance length over multiple interactions. For utterances to be reduced, of course, they must vary in length. Motivated by our empirical observation that meaningful clauses are the primary unit of reduction, we extend our grammar to include *conjunctions*. This is one of the simplest ways to constructing longer utterances composition-

ally from lexical primitives, using the product rule:

$$\mathcal{L}(u_i \text{ and } u_j, o) = \mathcal{L}(u_i, o) \times \mathcal{L}(u_j, o)$$

Analogous to our tangram stimuli, which have many ambiguous features and figurative perspectives that may be evoked in speaker descriptions, we consider a simplified scenario where speakers can refer to two different features of the two objects $\{o_1, o_2\}$. The speaker has four primitive words at their disposal – two words for shape ($\{u_{s1}, u_{s2}\}$) and two for color ($\{u_{c1}, u_{c2}\}$) – and has uncertainty over the initial meanings of all four.

While we established in the previous section that conventions can emerge over a reference game in the complete absence of initial preferences, players often bring such preferences to the table. A player who hears ‘ice skater’ on the first round of our tangrams task is more likely to select some objects more than others, even though they still have some uncertainty over its meaning in the context. To show that our model can accommodate this fact, we allow the speaker’s initial prior meanings to be slightly biased. u_{s1} and u_{c1} are more likely to mean o_1 ; u_{s2} and u_{c2} are more likely to mean o_2 .

We ran 1000 forward samples of 6 rounds of speaker-listener interaction, and averaged over the utterance length at each round.⁶ Our results are shown in Figure 4C: the expected utterance length decreases systematically over each round. To illustrate in more detail how this dynamic is driven by an initial rational preference for redundancy relaxing as reference becomes more reliable, we walk step-by-step through a single trajectory.

Consider a speaker who wants to refer to object o_1 . They believe their knowledgeable partner is slightly more likely to interpret their language using a lexicon in which u_{s1} and u_{c1} apply to this object, due to their initial bias. However, there is still a reasonable chance that one or the other alone actually refers strongly to o_2 in the true lexicon. Thus, it is useful to produce the conjunction “ u_{s1} and u_{c1} ” to hedge against this possibility, despite its higher cost. Upon observing the

⁶In our simulations, we used $\alpha = 10$ and found the basic reduction effect over a range of different biases

listener’s response (say, o_1), the evidence is indeterminate about the separate meanings of u_{s1} and u_{c1} but both become increasingly likely to refer to o_1 . In the trade-off between informativity and cost, the shorter utterances remain probable options. Once the speaker chooses one of them, the symmetry collapses and that utterance remains most probable in future rounds. In this way, meaningful sub-phrases are omitted over time as the speaker becomes more confident about the true lexicon.

General Discussion

In this paper, we revisited the classic phenomenon of convention-formation in a large-scale, text-based replication of the tangrams task. We argued that several key qualitative patterns in the data – arbitrariness, stability, and the reduction of utterance length over repeated interactions – can be explained by our model of informative communication under lexical uncertainty. This model formalizes a theory where conventions emerge via uncertain agents who assume their partner is knowledgeable and informatively using some conventional lexicon. Through repeated observations of their partner’s actions, agents learn this lexicon, thereby coordinating and aligning to one another.

Theories of convention-formation vary in the extent to which social reasoning about common ground is required. Our agents lie on a spectrum between the heuristic updating agents of Barr (2004) and the sophisticated agents of Clark & Wilkes-Gibbs (1986), who collaboratively build up explicit representations of mutual knowledge. Speakers and listeners in our model implicitly coordinate their beliefs through a shared history of observations, which serves as “common ground” in an informal sense. They make critical use of pragmatic, social reasoning in order to learn meanings, but do not explicitly consider the fact that this history is shared, or represent their partner’s own uncertainty.

By capturing reduction, which purely heuristic theories have not yet demonstrated, we showed that minimal assumptions of social reasoning go a long way in accounting for key phenomena. Still, our model falls short in some ways. For instance, because we do not provide a mechanisms for the listener agent to respond with confirmation, repair, or follow-up questions, we cannot make explicit predictions about the reduction in *listener messages* (as shown in Fig. 2) or the effect of listener input on the conventionalization process. These phenomena require our model to deal with planning over extended dialogues, and more sophisticated speech acts. Similarly, while our model was explicitly designed with linguistic conventions in mind, it remains to be seen whether the same formulation generalizes to broader behavioral conventions. For example, the real-time coordination games used in Hawkins & Goldstone (2016) may not require players to reason about a structured lexicon with noise, but an action policy representation may play a similar role. While there remain many complex aspects of convention-formation in communication games left for future research, our approach nonethe-

less serves as a lower bound on the degree of social reasoning needed to capture lexical conventions in these games.

Acknowledgements

We thank Nicole Maslan for her contributions during piloting. This work was supported by ONR grant N00014-13-1-0788 and a Sloan Research Fellowship to NDG. RXDH was supported by the Stanford Graduate Fellowship and the National Science Foundation Graduate Research Fellowship under Grant No. DGE-114747.

References

- Barr, D. J. (2004). Establishing conventional communication systems: Is common knowledge necessary? *Cognitive Science*, 28(6), 937–962.
- Bergen, L., Levy, R., & Goodman, N. D. (2016). Pragmatic reasoning through semantic inference. *Semantics and Pragmatics*, 9(20).
- Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(6), 1482.
- Centola, D., & Baronchelli, A. (2015). The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proceedings of the National Academy of Sciences*, 112(7), 1989–1994.
- Clark, H. H., & Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition*, 22(1), 1–39.
- Frank, M. C., Goodman, N. D., & Tenenbaum, J. B. (2009). Using speakers’ referential intentions to model early cross-situational word learning. *Psychological Science*, 20(5), 578–585.
- Garrod, S., & Doherty, G. (1994). Conversation, co-ordination and convention: An empirical investigation of how groups establish linguistic conventions. *Cognition*, 53(3), 181–215.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11), 818–829.
- Goodman, N. D., & Stuhlmiller, A. (2013). Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science*, 5(1), 173–184.
- Goodman, N. D., & Stuhlmiller, A. (electronic). The design and implementation of probabilistic programming languages. Retrieved from <http://dippl.org>
- Graf, C., Degen, J., 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*.
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and semantics* (pp. 43–58). New York: Academic Press.
- Hawkins, R. X. D. (2015). Conducting real-time multiplayer experiments on the web. *Behavior Research Methods*, 47(4), 966–976.
- Hawkins, R. X. D., & Goldstone, R. L. (2016). The formation of social conventions in real-time environments. *PLoS ONE*, 11(3), 1–14.
- Krauss, R. M., & Weinheimer, S. (1964). Changes in reference phrases as a function of frequency of usage in social interaction: A preliminary study. *Psychonomic Science*, 1(1-12), 113–114.
- Lewis, D. (1969). *Convention: A philosophical study*. Harvard University Press.
- Schuster, S., & Manning, C. D. (2016). Enhanced english universal dependencies: An improved representation for natural language understanding tasks. In *LREC 2016*.
- Smith, N. J., Goodman, N. D., & Frank, M. C. (2013). Learning and using language via recursive pragmatic reasoning about other agents. In *NIPS* (pp. 3039–3047).
- Toutanova, K., Klein, D., Manning, C. D., & Singer, Y. (2003). Feature-rich part-of-speech tagging with a cyclic dependency network. In *NAACL-HLT* (pp. 173–180).
- Wilkes-Gibbs, D., & Clark, H. H. (1992). Coordinating beliefs in conversation. *Journal of Memory and Language*, 31(2), 183–194.
- Young, H. P. (2015). The evolution of social norms. *Annual Review of Economics*, 7, 359–387.