Social Conventions in Collaborative Tagging

Christopher W. Nell

April 22, 2009

Abstract

With the increasing popularity of online collaborative tagging services, the study of user tagging behavior is of significant interest. A model of vocabulary evolution is described from the literature, re-interpreted using game theoretic concepts, and tested via simulation. Refinements are motivated in view of the nature of collaborative tagging and in light of recent analyses of real-world collaborative tagging data.

1 Introduction

The idea of applying game theory to the study of language evolution is not new. In introducing language games [8], Steels described a simple process by which an uncoordinated population of boundedly-rational agents reach the convention of assigning common terms to objects. While not formally presented in such terms, the process he described is similar to the stochastic coordination games studied in contemporary game theoretic work [6]. However, perhaps because the restricted interactions assumed in practical implementations do not closely resemble natural language, this work has not seen wide application to problems in computational linguistics.

Since the mid-2000s there has been a rise in the popularity of online collaborative tagging systems such as Delicious. * Users of these services annotate web resources using short keywords; the process of tagging a new resource can be seen as a naming process under a greatly simplified grammar. As such, there has been a resurgence of interest in the study of language evolution in the context of collaborative tagging. For example, a group of Italian physicists have been actively investigating the evolution of vocabulary (or *semiotic dynamics*) in collaborative tagging systems. Baronchelli *et al.* motivate a simplified language game as being relevant to collaborative tagging and discuss its dynamics, but do not compare to real tagging data or incorporate any domain-specific considerations into the model [1]. Meanwhile, Cattuto *et al.* develop a stochastic model which explains phenomena observed in real tagging data, but do not explicitly connect their work to game theory [2].

*http://www.delicious.com

This report presents a survey of recent work on agent-based simulation of semiotic dynamics, with a focus on modeling the collaborative tagging process. In Section 2, a game theoretic formalism describing emergent social conventions is introduced. In Section 3, the Naming Game is described and re-interpreted using these game-theoretic formalisms, and tested via simulation. It is further refined to model collaborative tagging in Section 4, through examining the user tagging process and incorporation of results from recent analyses of tag data. Finally, in Section 5 the results of these investigations are summarized, and avenues for further research are identified.

Throughout the text, the reader is assumed familiar with fundamental concepts and notations of game theory, set theory, and statistics.

2 Social conventions in stochastic games

In their 1997 paper, Shoham and Tennenholtz [6] introduce a game theoretic framework for describing the emergence of social conventions and apply it to a restricted class of games. Their terminology is adopted here, and important results are summarized. Definition 1 below is reproduced verbatim:

Definition 1 (n-k-g stochastic social game)
An n-k-g stochastic social game consists of a set of n agents, a k-person game g, and an unbounded sequence of ordered tuples of k agents selected from a uniform distribution over the n given agents.

Repeated n-k-g stochastic social games can be studied in terms of the common *selection function* used by agents to choose actions in each stage game. It is assumed that there are no extrinsic preferences for particular actions based on information not encoded into the game itself. Of special interest are settings in which only limited, *local information* is available to the an agent; here local information means restricting to the agent's past action choices and realized payoffs.

Highest Cumulative Reward (HCR), in which in which agents simply play the action that has recently been most profitable, is an example local selection function introduced by Shoham and Tennenholtz. Despite the fact that it does not allow agents to reason at all, they establish a variety of results for HCR in the case of n-2-g stochastic social games, where g is of a restricted class of symmetric 2×2 normal-form games that includes both coordination and cooperation (prisoner's dilemma). Ultimately, they show that all n agents will eventually converge to a stable, rational, pure strategy profile.[†] When this occurs, a *social convention* is said to have emerged.

Definition 2 (social convention)
A social convention is a restriction of all agents to a single strategy profile.

[†] "Rational" means each agent's payoff is strictly grater than the associated maxmin value.

The key message is that – in this restricted setting – it is possible even for information-limited, extremely boundedly-rational agents to reach social conventions without external coordination. This will not necessarily happen quickly; indeed, a lower-bound of $\Omega(n \log(n))$ game iterations is proven necessary for any local selection function to produce a social convention. Through simulation it is established that the specific dynamics with which a social convention is reached depend strongly on the particular game played.

In the broader context of games with more than two actions (such as the games discussed in this paper), a different selection function is introduced. In *External Majority* (EM), each agent plays the strategy most commonly observed in other agents. This rule is called *semi-local* because in addition to local information, the agent has access to information about directly-observed second-party actions. No theoretical results are provided, but a simulation study of EM on a class of *extended coordination games* suggests that increasing the number of actions only sub-logarithmically decreases the probability of a social convention emerging.

3 Modeling language evolution

Although the underlying mechanisms for modeling language evolution have often been loosely described as games, they are typically not presented in a manner consistent with game theoretic conventions. Moreover, there has not been an effort to update these basic game models in light of the statistical and graph-theoretic analyses presented in recent studies of semiotic dynamics in social tagging. As such, the contributions of this work are twofold – first, to provide formal game-theoretic models of vocabulary evolution; and second, to incorporate collaborative tagging domain observations into these games. The first of these goals is addressed in this section.

3.1 A procedural view of language games

Steels [8] introduces the concept of *language games* with a high-level description of the interaction process, and in particular does not adopt a game theoretic perspective. Despite this, a restricted version of the process he describes can be understood as a stochastic social game. As such, language games are first described here in procedural terms (although using more precise notation), before being simplified and mapped into the framework of Section 2.

Definition 3 (language game setting)
A language game setting consists of sets of agents N, objects O, features F, and words V. Each $i \in N$ has three associated private mappings – a word mapping *$U_i : V \to \wp(F)$, a* feature mapping *$C_i : F \to \wp(V)$, and a* success mapping *$P_i : V \times F \to \mathbb{Z}^2$ – and two distributions over V – ψ_i and ϕ_i. Initially $U_i(v) = \emptyset$, $C_i(f) = \emptyset$, and $P_i(v, f) = (0, 0) \quad \forall\, (v, f) \in V \times F$. Each $o \in O$ has an associated set of identifying features $F_o \subseteq F$.*

An iteration of communication in the particular language game described by Steels proceeds as follows:

1. A pair of agents s and h, and one object o, are chosen uniformly at random.
2. Agent s, the *speaker*, selects a subset of the object features $F_o^s \subseteq F_o$, and updates word and feature mappings as necessary according to:

$$\forall f \in \{f' \in F_o^s \mid C_s(f') = \emptyset\} : \\ \left.\begin{array}{rl} C_s(f) & \equiv \{v_f\} \\ U_s(v_f) & \equiv \{f\} \end{array}\right\} \quad \text{where } v_f \sim \psi_s(v \mid U_s)$$

3. Starting from $V_o^s = \emptyset$, the speaker constructs a signal to send via the following selection function:

$$\begin{aligned} &\forall f \in F_o^s : \\ &\qquad V_o^s = V_o^s \cup \{v_f\} \quad \text{where } v_f \sim \phi_s(v \mid C_s, P_s, f) \\ &P_s(v_f, f) = P_s(v_f, f) + (1, 0) \end{aligned}$$

4. Agent h, the *hearer*, receives this signal; the successfully-communicated words and associated meanings from the perspective of each agent are:

$$\begin{aligned} V_o^{sh} &\equiv \{v \in V_o^s \mid F_o \cap U_h(v) \neq \emptyset\} \\ F_o^{sh} &\equiv \{f \in F_o \mid V_0^{sh} \cap C_h(f) \neq \emptyset\} \\ X_o^s &\equiv \{(v, f) \in V_o^{sh} \times F_o \mid f \in U_s(v)\} \\ X_o^h &\equiv \{(v, f) \in V_o^{sh} \times F_o \mid f \in U_h(v)\} \end{aligned}$$

5. The agents' success mappings are updated via:

$$\begin{aligned} P_s(v, f) &= P_s(v, f) + (0, 1) \quad \forall (v, f) \in X_o^s \\ P_h(v, f) &= P_h(v, f) + (1, 1) \quad \forall (v, f) \in X_o^h \end{aligned}$$

6. If $V_o^{sh} = V_o^s$ the communication is deemed a success. Otherwise, the hearer's word and feature mappings are updated:

$$\begin{aligned} U_h(v) &= U_h(v) \cup (F_o \backslash F_o^{sh}) \quad \forall v \in V_o^s \backslash V_o^{sh} \\ C_h(f) &= C_h(f) \cup (V_o^s \backslash V_o^{sh}) \quad \forall f \in F_o \backslash F_o^{sh} \end{aligned}$$

Simulation results suggest that convergence to a social convention occurs in such language games, and can do so quickly – within a few thousand iterations in the case of 5 agents and 16 objects. Additionally, the model admits many interesting phenomena observed in real language; for example, both synonymy and homonymy are observed to arise. Unfortunately, the model also has significant limitations. First, the game is sufficiently complex that no theoretical analysis is presented, and is too complex to be directly expressed under the framework of Section 2. Moreover, the model assumes agents have unlimited memory – not only do they remember every word-feature association ever encountered, but they are unable to forget any inappropriate pairings that may have been made. Finally, communication is largely unidirectional – the speaker receives no information other than indications of comprehension from the listener.

3.2 The Naming Game

In order to address the practical issues with language games, various simplifications have been proposed. One relatively recent example is that of Baronchelli *et al.* [1]. In their work they consider a significantly restricted language game instance called the *Naming Game*. In terms of Definition 3, the restrictions are:

$$
\begin{aligned}
O &= \{o\} && \text{(a single object)} \\
F = F_o &= \{f\} && \text{(a single feature)} \\
|V| &= \aleph_0 && \text{(an unlimited vocabulary)} \\
\psi_i(v \mid U_i) &= \begin{cases} 1 & \text{if } U_i(v) = \emptyset \\ 0 & \text{otherwise} \end{cases} && \text{(words invented uniformly at random)} \\
\phi_i(v \mid C_i, f) &= \begin{cases} 1 & \text{if } v \in C_i(f) \\ 0 & \text{otherwise} \end{cases} && \text{(signal chosen arbitrarily from options)}
\end{aligned}
$$

Finally, memory is limited by having agents forget all previous word-feature associations after each successful communication. Specifically, step 6 of Section 3.1 is replaced with the following (in terms of the above simplifications rather than the general case):

6. If $V_o^{sh} = V_o^s$ the communication is deemed a success, and both agents' feature mappings are updated:

$$
\begin{aligned}
C_s(f) &= V_o^s \\
C_h(f) &= V_o^s
\end{aligned}
$$

   Otherwise, the hearer's feature mappings are updated:

$$
C_h(f) = C_h(f) \cup V_o^s
$$

These restrictions have the following important implications on the game:

- all signals consist of a single word
- homonymy cannot be modeled as only one feature exists
- at most $|N|/2$ unique words are created, one by each agent whose first interaction is as a speaker
- success mappings P_a are never used, so are omitted

Baronchelli *et al.* demonstrate via simulation that this simple process produces common-terminology social conventions. Moreover, the adoption of convention is observed to occur in a relatively sudden transition after a long period of communication using a large vocabulary. This qualitative behavior is very similar to that observed in natural language [5], and in particular is observed to be consistent with that expected of an urn model [3]. While analytic explanations for these observations are presented, bounds on the rate of convergence are not proved. Finally, no alternative models are tested alongside the Naming Game.

3.3 Game theoretic representation of language evolution

Although not introduced in such terms, the Naming Game can be modeled as an n-2-g stochastic social game with $n = |N|$ agents, where g is a variation of the coordination game called the *comprehension game* (to avoid confusion with the non-game-theoretic language game) with $s = n/2$:

Definition 4 (comprehension game)
A comprehension game is an asymmetric 2-person game, where player 1 plays $a_1 \in A_1$ where $|A_1| = s$, and player 2 plays $a_2 \in A_2 \equiv \wp(A_1)$ so that $|A_2| = 2^s$. The payoff for both agents is $x > 0$ if and only if $a_1 \in a_2$, and it is $-x$ otherwise.

While this game is uninteresting if we assume rational agents with complete knowledge of the game (player 2 should simply play the action corresponding to the full set of s actions), interesting behavior can arise if action choices are constrained by a selection function. The mechanism of Section 3.2 implicitly defines a semi-local selection function for a comprehension game, here called *Reward Restarted Uniform* (RRU):

Definition 5 (RRU for comprehension games)
Under RRU, agent i as player 1 in a comprehension game selects a uniformly-mixed strategy s_1^i over all previously-observed actions $A_1^i \subseteq A_1$ (or all actions A_1 if $A_1^i = \emptyset$). Agent j as player 2 selects the pure strategy s_2^j corresponding to $A_1^j \in A_2$. On positive payoff, both agents' memory of all but the realized action a_1 is reset; i.e. $A_1^i = A_1^j = \{a_1\}$.

While this selection function reproduces the Naming Game with $s =$ player 1 and $h =$ player 2, it is not the only possible choice in this setting.[‡] In particular, both the HCR and EM can be adapted to comprehension games:

Definition 6 (HCR for comprehension games)
Under HCR, agent i as player 1 in a comprehension game selects an action $a_1 \in A_1$ associated with the highest cumulative payoff, where only payoffs associated with agent i's recent actions as player 1, ${A'_1}^i \subseteq A_1$, are counted.

Definition 7 (EM for comprehension games)
Under EM, agent i as player 1 in a comprehension game selects the action

[‡] The strategy of player 2 cannot be changed without sacrificing the intended interpretation of the comprehension game, but the strategy of player 1 is not so restricted.

$a_1 \in A_1$ which is observed the most often in other agents, where all recently observed actions $A_1^i \subseteq A_1$ are counted.

In either case, limited memory can be enforced by bounding the size of A_1^i or ${A'_1}^i$. The main difference between HCR and EM is that agents condition on actions taken when player 1 under HCR, and on observations made when player 2 under EM. The two sources are combined under *semi-local HCR*:

Definition 8 (semi-local HCR for comprehension games)
Under semi-local HCR, agent i as player 1 in a comprehension game selects an action $a_1 \in A_1$ associated with the highest cumulative payoff, where all recently observed actions $A_1^i \subseteq A_1$ are counted.

In this factored representation, it is not clear whether the observations of Baronchelli *et al.* [1] noted in Section 3.2 are specific to RRU, or are typical of the collaboration game in general. In order to probe this issue, simulations of comprehension games using both RRU and memory limited semi-local HCR were conducted; the results are summarized in Figure 1. This figure is qualitatively identical to the reports reported by Baronchelli *et al*, indicating that RRU is indeed equivalent to the Naming Game. From this figure it is also clear that convention emergence occurs for both selection functions, and indeed faster under semi-local HCR than RRU. This strongly suggests that the underlying comprehension game, and not the specific selection function, is fundamentally responsible for these dynamics.

3.4 Beyond single features

Perhaps the most restrictive limitation in modeling language evolution by an n-2-g stochastic social comprehension game is that only one feature per object is supported – a carryover from the Naming Game. The original language game definition does not have this restriction, and it would be useful to provide a similarly unrestricted game theoretic model. This might be achieved by extending current framework to a Bayesian game setting – the speaker's type determining which feature a is communicated, and the hearer's type determining which feature the signal is interpreted as corresponding to. Clearly, this approach requires substantial development before being practically useful.

4 Modeling collaborative tagging

Having formally described the Naming Game as a n-2-g stochastic social comprehension game with a specific choice of selection function, focus shifts to the appropriateness of this model for describing collaborative tagging. In this section, collaborative tagging is formally described and modeled using language games. Modifications to the n-2-g stochastic social comprehension game model are proposed, and a specific selection function is motivated in light of the results of recent analyses of tag data.

4.1 A procedural view of collaborative tagging

Communication via a collaborative tagging service t regarding a web resource o with true features F_o typically follows the following general procedure, using the terminology of Section 3.1: [§]

1. Service t sends signal V_o^t summarizing features $F_o^t \subseteq F_o$, using some function ϕ_t
2. User i maps V_o^t to a set of features F_o^{ti}
3. User i observes o and identifies a set of relevant features F_o^i
4. User i sends signal V_o^i summarizing $F_o^i \subseteq F_o$, using some function ϕ_i which may depend on F_o^{ti}
5. Service t maps V_o^i to a set of features ${F'}_o^i$
6. Service t updates F_o^t in light of ${F'}_o^i$

In this notation, each signal V_o is a set of tags, and the resource features F_o are the underlying concepts described by the tags. It thus appears that in principle tagging can be described by a language game involving pairwise, bidirectional communication between users and the service. However, a fundamental difference between this process and all models described so far is that it is *centered* – a unique center agent t is involved in each interaction.

4.2 The Naming Game for collaborative tagging

Baronchelli *et al.* [1] specifically introduce their work as a model appropriate for understanding trends in collaborative tagging. However, they do not explicitly describe how their model relates to a typical tagging process such as the one described here.

As has been discussed, the Naming Game treats the vocabulary used to describe different objects as independent. It also assumes $F_o = \{f\}$, so documents with multiple features can only be modeled by assuming these features are independent. Finally, it assumes a unidirectional communication between arbitrarily chosen agents. This last point is fundamentally incompatible with the tagging process, which involves bidirectional communication between an arbitrary agent and the center. To address this incompatibility, a centered version of the n-k-g stochastic social game is introduced:

Definition 9 (t-n-k-g centered stochastic social game)
A t-n-k-g centered stochastic social game consists of a set of a central agent t, n non-center agents, a k-person game g, and an unbounded sequence of ordered tuples of t and $k-1$ agents selected from a uniform distribution over the n non-center agents.

[§] In some use cases, step 3 may occur first.

Collaborative tagging can then be seen as a t-n-2-g centered stochastic social game, where g is a *pairwise comprehension* game:

Definition 10 (pairwise comprehension game) *A pairwise comprehension game between agents i and j is a sequence of two comprehension games; the first with i as player 1 and j as player 2, and the second with roles reversed. Each agent receives payoff equal to the sum of payoffs in the two stage games.*

If both agents use RRU, this game is can be interpreted an application of the Naming Game to collaborative tagging. However, as indicated by the results of Section 3.3, RRU is not necessarily the best choice of selection function. To motivate alternatives, it is useful to consider analyses of actual tagging data.

4.3 Statistical models of tag data

In this paper, game theoretic models of emergent conventions corresponding to vocabulary evolution are described and developed. However, it is not necessary to model agent interactions in order to describe these equilibria. Indeed, several studies have instead focused on statistical explanations of patterns observed in actual collaborative tagging data. In a 2006 paper, Golder and Huberman [4] present an exploratory analysis of many aspects of the Delicious collaborative tagging service. Among other observations, they informally observe that the distribution of document tags is similar to what would be observed in an urn process [3]. Cattuto *et al.* [2] report a study with qualitatively similar results, but also explicitly develop and justify (both theoretically and experimentally) a probabilistic model of tagging. Their model is memory-limited version of the Yule-Simon process [7, 9], an urn model like that alluded to by Golder and Huberman. Their model inspires the *memory-limited Yule-Simon* (MYS) selection function for non-center agents:

Definition 11 (MYS for comprehension games)
Under MYS, agent i as player 1 in a comprehension game selects an action $a_1 \in A_1$ according to a memory-limited Yule-Simon process:

- *with probability p, a_1 is chosen randomly from A_1*
- *with probability $1-p$, a position x is sampled with probability proportional to $\frac{1}{x-\tau}$ for some parameter τ. The action $(A_1^i)_x \in A_1^i$ which was xth-most recently observed is selected, where all recently observed actions $A_1^i \subseteq A_1$ are considered.*

The $1-p$ case is motivated by the observation that more recent tags are (much) more likely to be seen and emulated by users. By fitting the parameters p and τ[¶], Cattuto *et al.* [2] observed a very good fit between the model's predictions and tag distributions observed on documents.

[¶]Typical values are reported to be $p = 0.03$ and $\tau = 40$

5 Conclusions

In this work, a simple "Naming Game" describing vocabulary evolution via uncoordinated agent interactions was reported, and formalized in game theoretic terms through the introduction of stochastic social comprehension games. In this formalism, it is clear that the agents' selection function is a free design parameter, and a prominent alternative from game theory literature was tested and shown to exhibit qualitatively similar – and arguably superior – behavior to the implicit original choice. Approaches for applying this framework to the modeling of collaborative tagging was discussed, and an alternative selection function for non-central agent behavior was motivated based on previous analyses of tagging data.

There are at least three intuitive directions for future work. The first involves development of elaborated game theoretic models for language evolution. One possible such direction, extending stochastic social comprehension games into a Bayesian game setting, was introduced as a way to enable representation of multi-featural objects – a particularly relevant goal in the context of collaborative tagging. However, this direction has not been deeply pursued.

The second possibility is to perform more direct analysis of collaborative tagging data. In this work, the results of external analyses were considered, but insight was limited to the conclusions drawn by their respective authors. By performing first-party analysis, a more detailed understanding of the domain could be acquired as necessary.

Finally, the third and most interesting direction entails further exploration of specific game theoretic models for collaborative tagging. For example, although a data-supported model for users tag choice was proposed, it is still unclear what constitutes a good strategy for the center. Given the users' action dependence on communications from the center, the practical implications of this choice are significant. In a sense, this direction subsumes the previous two; both more capable models and a better understanding of the target domain are required to make meaningful progress.

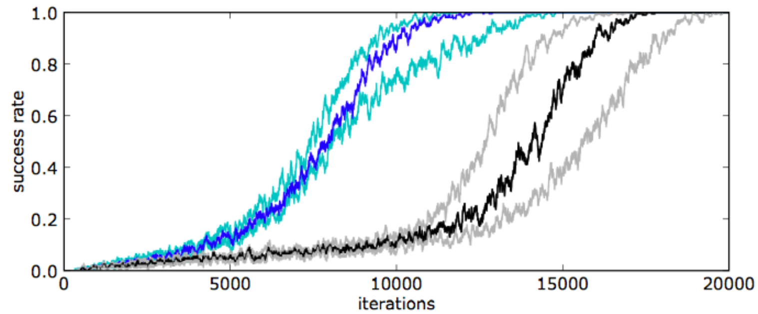

Figure 1: Communicative success rate evolution for a comprehension game with 500 agents. Curves represent median and quartile succcess rates over 10 simulations, measured in a 100-iteration sliding window. Blue curves correspond to 5-iteration-memory-limited semi-local HCR; black curves to RRU.

References

[1] Baronchelli, A., Felici, M., Caglioti, E., Loreto, V., & Steels, L. (2006). Sharp transition towards shared vocabularies in multi-agent systems. *Journal of Statistical Mechanics*, P06014.

[2] Cattuto, C., Loreto, V., & Pietronero, L. (2007). Semiotic dynamics and collaborative tagging. *Proc. National Academy of Sciences* 104, 1461.

[3] Eggenberger, F., & Polya, G. (1923). Uber die Statistik verketter vorgage. *Zeit. Angew. Math. Mech.* 1, 279–289.

[4] Golder, S., & Huberman, B. (2006). Usage patterns of collaborative tagging systems. *Journal of Information Science*, 32(2). 198–208.

[5] Lass, R. (1997). *Historical Linguistics and Language Change.* Cambridge University Press.

[6] Shoham, Y., & Tennenholtz, M. (1997). On the emergence of social conventions: Modeling, analysis and simulations. *Journal of Articial Intelligence* 94(1–2), 139–166.

[7] Simon, H. (1955). On a class of skew distribution functions. *Biometrika* 42, 425440.

[8] Steels, L. (1996). Emergent adaptive lexicons. In *Proc. Int'l Conf. on Simulation of Adaptive Behavior.*

[9] Yule, G. (1925). A Mathematical Theory of Evolution, based on the Conclusions of Dr. J. C. Willis, F.R.S. *Phil. Trans. of the Royal Soc. of London* B 213, 21-87.