Modeling the emergence of norms

(Extended Abstract)

Logan Brooks
University of Tulsa
Department of Computer Science
Tulsa, Oklahoma, USA
logan-brooks@utulsa.edu

Wayne Iba
Westmont College
Department of Computer Science
Santa Barbara, CA, USA
iba@westmont.edu

Sandip Sen
University of Tulsa
Department of Computer Science
Tulsa, Oklahoma, USA
sandip@utulsa.edu

ABSTRACT

Norms or conventions can be used as external correlating signals to promote coordination between rational agents and hence have merited in-depth study of the evolution and economics of norms both in the social sciences and in multiagent systems. While agent simulations can be used to gain a cursory idea of when and what norms can evolve, the estimations obtained by running simulations can be costly to obtain, provide no guarantees about the behavior of a system, and may overlook some rare occurrences. We use a theoretical approach to analyze a system of agents playing a convergence game and develop models that predict (a) how the system's behavior will change over time, (b) how much time it will take for it to converge to a stable state, and (c) how often the system will converge to a particular norm.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Measurement, Performance, Verification

Keywords

norm emergence, convergence

1. INTRODUCTION

The systematic study and development of robust mechanisms that facilitate emergence of stable, efficient norms via learning in agent societies promises to be a productive research area that can improve coordination in agent societies. Correspondingly, there has been a number of recent, mostly empirical, investigations in the multiagent systems literature on norm evolution under different assumptions about agent interaction frameworks, society topology, and observation capabilities [1, 2]. There is an associated need to develop analytical frameworks that can predict the trajectory of emergence and convergence of society-wide behaviors. Toward this end, we mathematically model the emergence of norms in societies of agents who adapt their likelihood of choosing one of a finite set of options based on their experience from repeated one-on-one interactions with other members in the society. The goal is to study both the process of emergence of norms as well as predict the likely final convention that is going to emerge if agents had preconceived biases or inclinations for certain options. We develop two different mathematical models under different interaction assumptions and validate model predictions using extensive simulations.

Cite as: Modeling the emergence of norms (Extended Abstract), Logan Brooks, Wayne Iba and Sandip Sen, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. 1239-1240.

2. PREDICTING NORM EMERGENCE

Consider a population of agents faced with a scenario where an agent interacts with exactly one other agent and each selects one of two actions (for example, driving on the right side of the road or the left). The goal for the agents is to interact in a coordinated manner; based on the outcome of their interaction (coordination or conflict), they adjust their predispositions to their selected actions.

In our models, an agent consists solely of a single number, p_i, representing the bias or probability of selecting one particular action. Agents select the other action with the complementary probability, $(1 - p_i)$. In our first model, every agent interacts with one other agent on every time step via $n/2$ random pairings for a population of n agents.

Based on the outcome of the interaction, the agent's bias is updated according to an update rule: $p_i(t+1) = p_i(t) \pm x$, where x, $0 < x < 1.0$, may be thought of as the learning rate and is typically small (*e.g.*, 0.01). This constant update is added so as to increase the likelihood of the action just chosen when it led to coordination, and is subtracted to decrease the action likelihood when it led to a conflict.

2.1 Full pairwise interaction

The expected fraction of agents from a population that will be coordinating with one another can be computed as $C = \frac{1}{n}\sum_{i=1}^{n} c_i$, where c_i is the probability that an agent i coordinates. In turn, we can define $c_i = p_i\bar{p}_{\hat{i}} + (1-p_i)(1-\bar{p}_{\hat{i}})$, where p_i is the probability agent i drives on the right and $\bar{p}_{\hat{i}}$ is the corresponding average likelihood across the population after removing the contribution of p_i from the population's average, $\bar{p}$. Note, $\bar{p}_{\hat{i}}$ can be calculated as $\bar{p}_{\hat{i}} = \frac{n \cdot \bar{p} - p_i}{n-1}$.

We can solve the recurrence relation for the mean bias in

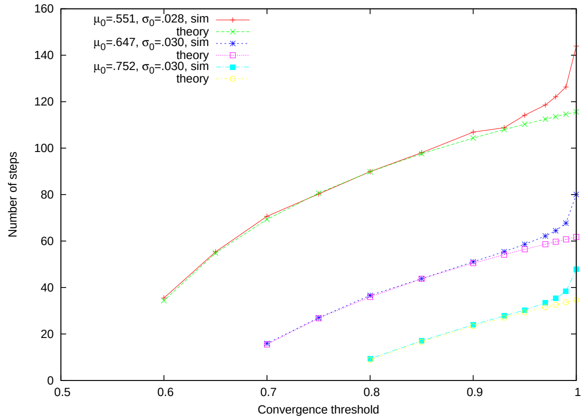

Figure 1: Comparing the number of time steps needed in the simulator and as given by the analysis as a function of initial population mean and convergence target.

the population at time t as follows:

$$\begin{aligned}\bar{p}(t) &= \bar{p}(t-1) + 2x\bar{p}(t-1) - x \\ &= (2x+1)\bar{p}(t-1) - x \\ &= y^t\bar{p}(0) - x\sum_{i=0}^{t-1} y^i \\ &= y^t\left(\bar{p}(0) - \frac{1}{2}\right) + \frac{1}{2}\end{aligned}$$

where $y = (2x+1)$. Since we want to know the number of time steps until the population settles on either driving on the right or the left, let us solve the above expression for t. By ignoring the $\frac{1}{2}$ that is added at the end of the expression we translate our interest from the range [0,1] to [-.5,+.5]. If we let $s = p(0) - \frac{1}{2}$, we want to see when the translated value exceeds 0.5 (or -0.5 for $p(0) < 0.5$)). If we allow some tolerance, ϵ ($\epsilon > 0$), then we care how the expression above relates to some limit, l^+, where $l^+ = \frac{1}{2} - \epsilon$ for populations converging to 0.5 in our translated frame of reference:

$$l^+ \leq y^t \cdot s \rightarrow t \geq \log_y l^+ - \log_y s.$$

For validating the theoretical predictions, we ran 50 simulations each with three populations of 100 agents each with initial bias means of 0.55, 0.65 and 0.75 (with $x = 0.01$). Figure 1 shows the number of time steps required as a function of a convergence threshold. Inspection of the figure indicates that the model accurately describes our empirical observations up to a convergence threshold of about 0.9.

2.2 Two-agent interaction

In the second model, we take a finer-grained look at the norm emergence process by selecting two agents, a_i and a_j, to interact on any given time step. The selected agents each calculate a random real number $r_k^t (k \in \{i,j\})$ from $U[0,1]$. Based on these random numbers, they each choose an action value $act_k^t = \begin{cases} +1, & r_k^t < p_k^t \\ -1, & r_k^t \geq p_k^t \end{cases}$. An action value of $+1$ indicates that the agent will choose to drive on the right side of the road, while a value of -1 corresponds to driving on the left side. If their actions did not coordinate, then each agent reduces the frequency with which it plays its chosen action. Mathematically, this can be expressed by:

$$p_k^{t+1} = \max\{0, \min\{1, p_k^t + act_i^t \cdot act_j^t \cdot \Delta_k^t(act_k)\}\},$$

where $\Delta_k^t(act_k) = x \cdot act_k$.

If $1/x$ is an integer and an agent is initialized with a p value that is a multiple of x, then we find that that agent's p value will always be a multiple of x. If there are n agents with p values constrained this way, then the population average, $\bar{p}$, can only assume values that are multiples of $\frac{x}{n}$, or $\frac{n}{x} + 1$ distinct values.

We can write an expression predicting the average convergence time and value for a given $\bar{p}$ value. Let $P(\bar{p})$ represent the estimated average convergence value for any population with an average bias of $\bar{p}$, and $T(\bar{p})$ be the expected number of time steps before converging. As with our treatment of the full pair-wise interaction, we ignore the corrections for values that fall below 0 or above 1. Consequently, we can express the value of $P(\bar{p})$ as a weighted average of the P values for all distributions that could be reached at the next time step. A similar expression can be used for $T(\bar{p})$, with an additional term of 1 to represent the current time step.

$$P(\bar{p}) = (1-\bar{p})^2 P\left(\bar{p} - 2\frac{x}{n}\right) + 2(1-\bar{p})\bar{p}P(\bar{p}) + \bar{p}^2 P\left(\bar{p} + 2\frac{x}{n}\right),$$

$$T(\bar{p}) = 1 + (1-\bar{p})^2 T\left(\bar{p} - 2\frac{x}{n}\right) + 2(1-\bar{p})\bar{p}T(\bar{p}) + \bar{p}^2 T\left(\bar{p} + 2\frac{x}{n}\right).$$

However, some values of P and T must be given in order to solve the system. Since $\bar{p}$ values of 0 or 1 indicate that the population has converged, we have definite values of P and T at these points: $P(0) = 0$, $P(1) = 1$, $T(0) = T(1) = 0$. The above equations for P and T form a nearly-diagonal linear system of equations, which can be solved in $O(n/x)$ time and space due to the discretization of the sample space. Solving this system of equations results in a close approximation of the average convergence time and values obtained in the simulations.

The predictions of the model were compared to the results of simulations in which all agents were initialized with identical p values. Due to space considerations, the results of this empirical evaluation are not shown here. However, for any starting $\bar{p}$ value, we found that the model very closely matched the simulation results for both average convergence value and time.

Between the two analyses presented in this paper, we establish a broad foundation for several types of subsequent work. For both analyses, we would like a better theoretical handle on how increasing diversity in the population impacts convergence time. In a similar vein, a more expansive analysis would provide insight into the effects that skewness in the population has on convergence.

3. REFERENCES

[1] J. Delgado, J. M. Pujol, and R. Sanguesa. Emergence of coordination in scale-free networks. *Web Intelligence and Agent Systems*, 1:131–138, 2003.

[2] S. Sen and S. Airiau. Emergence of norms through social learning. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, pages 1507–1512, 2007.