

Social Instruments for Robust Convention Emergence

Daniel Villatoro

IIIA-CSIC

Bellaterra, Barcelona, Spain

dvillatoro@iiia.csic.es

Jordi Sabater-Mir

IIIA-CSIC

Bellaterra, Barcelona, Spain

jsabater@iiia.csic.es

Sandip Sen

Department of Math. & Comp. Science

University of Tulsa, Oklahoma, USA

sandip-sen@utulsa.edu

Abstract

We present the notion of Social Instruments as mechanisms that facilitate the emergence of conventions from repeated interactions between members of a society. Specifically, we focus on two social instruments: *rewiring* and *observation*. Our main goal is to provide agents with tools that allow them to leverage their social network of interactions when effectively addressing coordination and learning problems, paying special attention to dissolving metastable subconventions. Our initial experiments throw some light on how Self-Reinforcing Substructures (SRS) in the network prevent full convergence to society-wide conventions, resulting in reduced convergence rates. The use of an effective composed social instrument, observation + rewiring, allow agents to achieve convergence by eliminating the subconventions that otherwise remained meta-stable.

1 Introduction

The social topology that restricts agent interactions plays a crucial role on any emergent phenomena resulting from those interactions [Kittock, 1993; Delgado, 2002; Urbano *et al.*, 2009]. In the literature on emergent behavior in multiagent systems (MAS), one active topic is *convention or norm emergence* as a mechanism for sustaining social order, increasing the predictability of behavior in the society and develop the details of those unwritten laws (we use the terms *conventions* and *norms* interchangeably in this paper). As conventions help agents to choose a solution from a space of equally good solutions (as long as all agents use the same choice), the selection of a coordination protocol, communication language, or the selection of the task to be executed in a multitask scenario are pertinent applications of conventions in MAS. In *social learning* [Mukherjee *et al.*, 2008; Sen and Airiau, 2007] of norms, where each agent is learning concurrently over repeated interactions with randomly selected neighbors in the social network, a key factor influencing success of an individual is how it learns from the “appropriate” agents in their social network.

A number of researchers from several communities (multiagent systems, physics, or economy) have studied the prop-

erties of underlying topologies in the convention emergence process [Kittock, 1993; Shoham and Tennenholtz, 1997; Urbano *et al.*, 2009; Delgado, 2002; Delgado *et al.*, 2003]. Most of these research consider a convention as emerged when 90% of the population has converged to the same convention. 90% convergence, however, should not be considered a convention (at least a robust one), as a convention, by definition, needs to be shared by the complete population.

In this paper, we show how agents can experience meta-stable subconventions depending on their position in the of interaction topology. As identified by several authors [Epstein, 2000; Toivonen *et al.*, 2009; Villatoro *et al.*, 2009], meta-stable subconventions interfere with the speed of the emergence of more general conventions. The problem of subconventions is a critical bottleneck that can derail emergence of conventions in agent societies and mechanisms need to be developed that can alleviate this problem. Subconventions are conventions adopted by a subset of agents in a society who have converged to a different convention than that adopted by the majority of the population. Subconventions are facilitated by the topological configuration of their social network, e.g., isolated areas of the graph promote endogamy, or by the agent reward function, e.g., concordance with previous history, promoting cultural maintenance, etc. Though it might be possible for agents to modify their reward functions in certain scenarios, we consider the general case where agents do not have that capability. Therefore the problem of subconventions has to be solved through the topological reconfiguration of the social network of agents.

We define *Social Instruments* to be a set of tools available to agents to be used within a society to influence, directly or indirectly, the behavior of its members by exploiting the structure of the social network. Social instruments are used independently (an agent does not need any other agent to use a social instrument) and have an aggregated global effect (the more agents use the social instrument, the stronger the effect). Specifically we focus on two social instruments: *rewiring* and *observation*. *Rewiring* allows agents to control the links that relate them with other agents by replacing them intelligently. This direct control of the topology of the social network allows agents to control whom they interact with, resulting in increased reward without actually altering the reward function. On the other hand, *observation* allows agents to obtain partial information of the convention emergence process by

observing other agents in the neighborhood. The access to this information allows agents to consider extra information over what they receive from direct interactions. This observation process also has an impact on the reward of the agents by facilitating their convention emergence within the society.

The main contribution of this work is the usage of Social Instruments as mechanisms that speed-up and support full-convergence (100% of the population using the same convention) in a range of social networks¹. Additionally, in the last part of the paper we argue how a combined social instrument allows the dissolution of subconventions in Scale-Free networks.

2 Proposed Social Instruments

2.1 Rewiring: Intelligent Link Removal and Creation

The Rewiring social instrument allows an agent to remove non-beneficial links with other agents, replacing them with new ones. Agents decide to rewire a link after the number of unsuccessful interactions² with another agent crosses a certain *Tolerance* threshold. Agents also need to decide whom they want to establish the new link with. We have developed three different methods:

(1) *Random Rewiring*: Agents rewire to a randomly selected agent from the population.

(2) *Neighbour's Advice*: Agents rewire to an agent recommended by a neighbour.

(3) *Global Advice*: Agents rewire to an agent that is randomly selected by the system from those that have the same strategy.

Though there are some similarities of our work with that of Griffith's [Griffiths and Luck, 2010], there is also a crucial difference. They use an *offline* evolutionary approach, where the system evolves over generations. On the other hand, we use an *online* approach where agents can adapt and modify their social network at runtime. In addition, our rewiring methods do not access any agent's private information (used only in the Global Advice method which is used as a control) such as their actual reward values.

2.2 Observation

In a social learning scenario, allowing agents to observe the strategy of agents outside their circle of interaction can provide useful information to support the convention emergence process. However, there should be a trade-off between observing and interacting. To analyze the effects of observation, we will allow agents to observe, at certain timesteps, a subset of other agents' states in the population. Therefore, agents are assigned an *Observation Probability*. Moreover, agents need to know the number of agents they can observe (*Observation Limit*) and who they can observe (*Observation Method*). We propose three different observation methods:

¹Initial experiments, unfortunately, have shown that full-convergence cannot be achieved in few specific configurations, e.g., scale free networks.

²Unsuccessful interaction in our convention emergence scenario corresponds to being uncoordinated or not sharing the same convention for that interaction.

(1) *Random Observation*: Agents observe random agents from the society.

(2) *Local Observation*: Agents observe their immediate neighbours in the social network.

(3) *Random Focal Observation*: Agents select a random agent from the society and observe that agent and its direct neighbors.

After the observation process, the agent will choose the majority action taken by the observed agents and will reinforce it.

Despite the similarity, this instrument and *mimicking* ([Hales and Arteconi, 2006]) behave differently. With observation, agents only access information that has been previously made public by the observed agent (agent's last played strategy), while with *mimicking*, they access information that can be considered private (complete list of neighbours and decision strategy function).

3 Model

The social learning situation for norm emergence that we are interested in is that of learning to reach a social convention. We use the following definition of a social convention: *A social law (a restriction on the set of actions available to agents) that restricts agents' behavior to one particular action is called a social convention* [Shoham and Tennenholtz, 1997]. Norms are implicit in our social learning formulation: Agents do not have any internal representation of norms, only preference for one action (the one specified by the norm) over the others.

For the sake of generalization, our framework is built with the most accepted convention emergence model (used by [Delgado *et al.*, 2003; Kittock, 1993; Mukherjee *et al.*, 2008; Sen and Airiau, 2007; Shoham and Tennenholtz, 1997; Walker and Wooldridge, 1995]): agents converge to a convention through repeated bilateral interactions with other agents from the neighborhood in their social network. Any interaction between two agents is represented as a 2-person *m*-action game. At each time step, each agent is paired with another agent and independently decide their actions. This decision is made without observing the other agent's identity or strategy. In our approach a social convention will be reached once *all* agents are in the same state or consistently choose the same action (the actual state reached or action chosen is immaterial).

As in several other research in convention emergence ([Delgado *et al.*, 2003; Kittock, 1993]), the social network of agents are chosen from the following structures: (i) a *one-dimensional lattice* with connections between all neighbouring vertex pairs (regular network); and (ii) a *scale-free network*, whose node degree distribution asymptotically follows a power law (irregular network).

As in [Kittock, 1993; Shoham and Tennenholtz, 1997; Villatoro *et al.*, 2009], agents are endowed a limited memory of past interactions (same size for all agents). Agents save in their memory *when* an interaction occurred, the *action* chosen, and the *reward* obtained. This information is used differently depending on the type of strategy decision procedure adopted. Agents cannot observe the memory, current deci-

sion, or reward obtained by the other agent, and hence cannot calculate the payoff for an action before interacting with the opponent.

Once the interaction model is fixed, we test our social instruments using three well-known strategy selection rules: (1) *Best Response Rule (BRR)* [Mukherjee *et al.*, 2008; Sen and Airiau, 2007], (2) *Highest Cumulative Reward Rule (HCRR)* [Shoham and Tennenholtz, 1997; Kittock, 1993], and, (3) *Memory Based Rule (MBR)* [Villatoro *et al.*, 2009].

The HCRR deterministically specifies the action that each agent has to take in each interaction. On the other hand, for BRR and MBR, agents use the Q-Learning algorithm to estimate the worth of each action, with an exploration rate of 25%.

4 Experiments

To create a manageable space of experimentation, some simulation parameters have been fixed: a population of 100 agents, with memory of size 5 (for HCRR and MBR), are located in a social network with different topologies: a low clustered³ one dimensional lattice (lattice with Neighborhood Size = 10), a high clustered one dimensional lattice (lattice with Neighborhood Size = 30), and a scale free network. Agents are initialized with no preference between the available actions, i.e., they randomly choose actions with equal probabilities. Presented results are averaged over 25 simulation runs.

4.1 Effects of Rewiring

We have explored the search space of the Tolerance Thresholds for the three rewiring methods, observing how it affects the convergence time and the number of links rewired when convergence is reached with the different strategy selection rules.

Influence of Rewiring Methods

In general, the *Global Advice (GA)* rewiring method produces the best convergence time due to its centralized nature and access to global information. Nonetheless the decentralized methods, specially the *Neighbour's Advice (NA)* method, also show good performances. The NA method improves the *Random Rewiring (RR)* method as it more expediently resolves the subconventions that appear in the one-dimensional lattices during the convention emergence process. Agents in the frontier use the rewiring instrument as they cross the tolerance level faster than those not in the frontier. For this reason, the RR method relinks an agent with a more suitable agent with a probability of $\frac{1}{\text{NumberOfActions}}$. In contrast, the NA method relinks the agent with another one with the same preference if it is accessible. In case there is no other agent with the same preference to connect with, random rewiring will be applied, obtaining in the worst case scenario, the same results. These results are also applicable for the scale-free networks.

Influence of topology

Regarding the effects of the topology, we find that the convergence time is increased under the effects of rewiring when

³Clustering Coefficient is a measure of degree to which nodes in a graph tend to cluster together.

the neighborhood size is increased. On regular networks, the diameter of the network directly affects the convergence times and the number of components. This effect is due to the clustering coefficient of the network. Lattices with higher neighborhood sizes are less fragmented than those with more restricted neighborhoods. Therefore, when increasing the neighborhood size, the number of links between agents also increases, thereby increasing the clustering coefficient. Highly clustered societies are more resistant to rewiring, as the node that wants to use the rewiring would have to apply it to a higher number of nodes, and then, be rewired to the same number of nodes with the appropriate strategy.

Experimental results also show interesting properties with Scale Free networks: when using the NA rewiring method, the number of components is significantly increased. As explained previously, rewiring is applied when two agents surpass their tolerance of unsuccessful interactions and NA will relink to a similar neighboring node. Because of the clustering coefficient of the Scale-Free networks⁴, NA will produce the disconnection of subgraphs from the main graph.

We can conclude that rewiring performs better in low clustered societies, producing a stratified population which results in significant reduction in convergence time. In more clustered networks, the tolerance level has to be chosen carefully (depending on the experiment) to produce an effective technique for norm emergence.

4.2 Effects of Observation

In this section we analyze the effects of *observation* as a social instrument when used by agents. We test and compare the three different methods proposed, exploring the search space with a representative range of *Observation Probability* values. We fix the *Observation Limit* to 10 for the experiments.

Influence of Observance Methods

Comparing the results from the three Observation methods, we observe that the *Random (RO)* and the *Random Focal Observation (RFO)* methods are the most effective ones, and have very similar results, when compared with the *Local Observation (LO)* method. The reason for this phenomenon can be attributed to the frontier effect. When agents use the LO method, they observe their direct neighbours. If the observing agent is in the frontier area, then, this observation is pointless. However, observing different areas gives a better understanding of the state of the world, and hence the RO and the RFO methods perform better.

Influence of topology

For the BRR and MBR strategy selection rules, we have observed that the different Observation methods produce a more pronounced effect in societies with higher diameters. We notice that a small percentage of Observation drastically reduces convergence times. The reason for this effect can also be attributed to the frontier and the subconvention effect previously discussed. Subconventions emerge more readily when

⁴The clustering coefficient distribution decreases as the node degree increases.

the social network has a small diameter and the frontier region represents the unsettled area. These subconventions are more easily resolved at these frontiers by observation rather than by learning through interactions.

5 Combining Instruments: Solving the Frontier Effect

Our initial results together with the observations made by other authors [Epstein, 2000; Toivonen *et al.*, 2009; Villatoro *et al.*, 2009] convinced us that subconventions are key obstacles to the emergence of global conventions. These subconventions thrive because of the topological structure of the network where they emerge. To achieve the dissolution of subconventions, they need to be resolved in what we identified as the “frontier” region.

A subconvention in a regular network is not meta-stable, but unfortunately, slows down the process of emergence. On the other hand, in other network types, such as random or scale-free, subconventions seem to reach meta-stable states⁵. Consequently, we have defined *weak frontiers* as the ones that are not meta-stable in regular networks, and *strong frontiers* as the ones generated in irregular networks.

By combining the social instruments presented in Section 2, we have designed a composed instrument for resolving subconventions in the frontier in an effective and robust manner. This composed instrument allows agents to “observe” when they are in a frontier, and then, apply rewiring, with the intention of breaking subconventions. To effectively use this combined approach, agents must first recognize when they are located on a frontier. We have previously defined a frontier as consisting of the group of nodes in the subconvention that are neighbours to other nodes with a different convention and that are not in the frontier with any other group.

The most important characteristic that defines a frontier is the existence of a confrontation. Confrontation occurs when two agents in an interaction do not share the same convention⁶.

Before proceeding further, we will define three characteristics of agents with respect to their convention and topological position in the network. An agent is *in equilibrium* if it has the same number of neighbours in its own convention as in the other convention. An agent is a *weak* node if the number of neighbours in its own convention is lower than those in the other, and an agent is a *strong* node otherwise (if the number of neighbours in its own convention is greater than those in the other).

In regular networks, two confronted agents are in a frontier region iff: (1) At least one of the confronted agents is in an equilibrium position, and (2) all the neighbours of an in-equilibrium confronted agent are strong nodes.

⁵By experimentation, we have observed that around 99% of the generated scale-free networks do not converge (to full convergence) up to one million timesteps with any of the decision making functions used in this work.

⁶Not sharing the same convention, choosing a different action, or choosing a different state to be, are considered equivalent expressions for our purpose.

By taking a snapshot at the end of the simulation of the emergence of irregular networks (such as scale-free topologies), we can extract the regions of the network that remain meta-stable with a convention different than the general convention. After compiling and studying those structures, we identified an abstract substructure that we have defined as *Self Reinforcing Substructures (SRS)*. These substructures, given the appropriate configuration of agents’ preferences, do maintain subconventions. These abstract structures are of two types, the *Claw* and the *Caterpillar* (see examples in Fig. 1), and can be found as subnetworks of scale-free and random networks.

The *Claw* SRS is formed by connecting a node with a number of *hangers*⁷ connected to it smaller than the number of links with the rest of the network. In the situation where the hangers coordinate to the same convention among themselves and with the connecting node, we have a self-reinforcing structure. For example, in Fig. 1(a), A is the central node, having one connection with the rest of the network and 3 hangers: B (that it is another claw), C (plain hanger) and D (chain’s connecting node).

The *Caterpillar* SRS is a structure formed by a central path and from its members can hang other SRSs (such as claws, chains, or plain hangers). For example, in Fig. 1(b), A, B, C, and D are members of the central path, and the other nodes reinforce them.

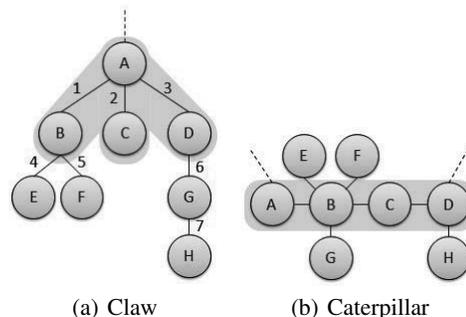


Figure 1: Self-Reinforcing Structures.

As we have observed, the existence of these SRSS (74% of the generated networks with the methods described in [Delgado *et al.*, 2003] contain SRS) are the main reason why convergence to a 90% level (as observed by [Delgado *et al.*, 2003]) is achieved relatively quickly, but overcoming the last 10% (containing the SRS) is much harder to achieve.

Inspired by a previous paper [Sen and Airiau, 2007], we have performed a test varying the amount of players with a fixed strategy to observe the dynamics of the emergence of conventions. As shown in Fig. 2, the emergence of conventions in scale-free networks follow the same behavioural pattern amongst them with different amount of fixed players,

⁷A *hanger* is formed by nodes that are connected to a member of a cyclic component, but which do not themselves lie on a cycle [Scott, 2000], and a *chain* is a walk in which all vertices and edges are distinct.

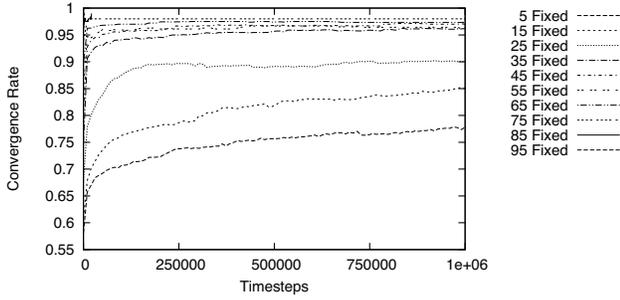


Figure 2: Evolution of Conventions in Scale-Free Networks with Fixed Players.

following a power-law distribution (as the node-degree distribution). Convention emergence is achieved relatively fast for the majority of the network. In the remaining cases where SRSs are located, however, a significant delay is observed.

By giving agents the instruments (an agent will use observation to realize it is part of a SRS, and rewiring to dissolve it) to dissolve these SRS, we hypothesize that convention emergence will be expedited and full convergence rates will be obtained.

5.1 Results with Combined Instruments

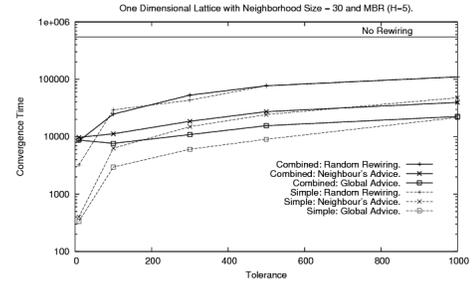
We have conducted exhaustive experimentation with the combined instrument on the three topologies and using the different decision making functions described in the previous section. The use of the composed instrument on the regular networks does not produce an improvement on convergence time with respect to simple rewiring (one specific scenario of topology and strategy decision technique can be observed in Fig. 3(a)).

However, an important improvement is observed in the number of rewired links (one example of this improvement can be seen in Fig. 3(b)). In general, this improvement is observed for lower tolerances. The reason for this effect is that for higher tolerances rewiring works in the same way as the composed social instrument, but without observing. For those smaller values, the effect is significant, reducing the number of rewiring links down to half of the original value.

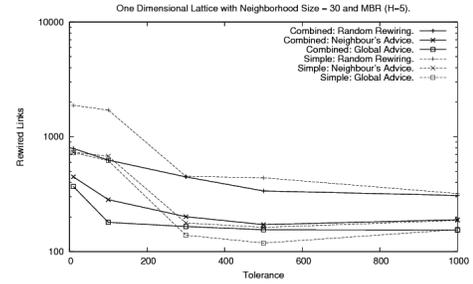
On the other hand we observe an important improvement for convergence times when using the composed instrument (with the recognition of SRS) on irregular networks. The results presented in Figure 4 represent the average results from 25 different scale-free networks with and without using the Combined Social instrument. By comparing Figure 4(a) and Figure 4(b) we notice the tradeoff between the improvement in convergence time and the amount of rewiring needed. The reason for this phenomena is because the Composed social instrument decomposes the SRS differently than simple rewiring which only rewires the node in the actual frontier.

6 Conclusions and Future Work

We have introduced the use of *Social Instruments* as tools that facilitate norm evolution. We have identified the characteristics and opportunities for effectively utilizing these social

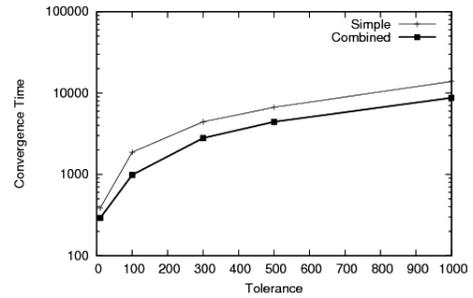


(a) Convergence Time

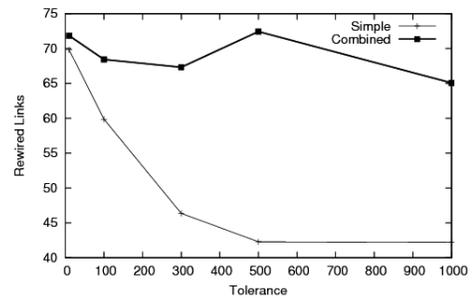


(b) Rewired Links

Figure 3: Comparison with Simple and Combined Social Instruments on Regular Network using MBR.



(a) Convergence Time



(b) Rewired Links

Figure 4: Comparison with Simple and Combined Social Instruments on Scale Free Network using BRR.

instruments for facilitating norm emergence through social learning. Social instruments are attractive since they do not require centralized monitoring or enforcement mechanisms, normally are extremely easy to use, have very low computational costs, and are scalable to large systems.

Experimental results with the identified social instruments have shown that the emergence of transitory subconventions are the cause of the delay of the emergence of global conventions. From results presented in this paper for the two simple social instruments studied, we observe that the most effective social instruments are those that more expediently solve this subconvention formation problem in the frontier regions.

To the best of our knowledge, this is the first attempt to achieve full convention emergence (100% of the population in the same state/choosing the same action) in scale-free networks. Other researchers studying norm emergence with these type of topologies chose a convergence rate of 90%, without considering that the rest 10% may not be possible to achieve [Kittock, 1993; Shoham and Tennenholtz, 1997; Urbano *et al.*, 2009; Delgado, 2002; Delgado *et al.*, 2003], due to the presence of the identified SRSs. We have presented a composed social instrument as a robust solution against the persistence of subconventions in social networks, improving the convergence times obtained with simple rewiring and finally achieving full convergence.

In a world where almost 950 million users belong to online social networking platforms (where virtual agents could also exist)⁸, it is important to understand what mechanisms these virtual entities should be equipped with to facilitate the emergence of common conventions (for the sake of the whole group) as quickly as possible. Moreover, as a system manager, the results from this work highlights the harmful potential of Self-Reinforcing Structures within the network for delaying the emergence process, and draws our attention to need for solutions to such critical problems.

Acknowledgments

This work was supported by the Spanish Education and Science Ministry [Engineering Self-* Virtually-Embedded Systems (EVE) project, TIN2009-14702-C02-01]; Mac-Norms Project [PIFCOO-08-00017] and the Generalitat de Catalunya [2005-SGR-00093]. Daniel Villatoro is supported by a CSIC predoctoral fellowship under JAE program. Sandip Sen is partially supported in part by an NSF award IIS-0757434. We also thank the CESGA and Rede Galega de Bioinformatica for the technical support.

References

[Delgado *et al.*, 2003] Jordi Delgado, Josep M. Pujol, and Ramón Sangüesa. Emergence of coordination in scale-free networks. *Web Intelli. and Agent Sys.*, 1(2):131–138, 2003.

- [Delgado, 2002] Jordi Delgado. Emergence of social conventions in complex networks. *Artificial Intelligence*, 141(1-2):171–185, October 2002.
- [Epstein, 2000] Joshua M. Epstein. Learning to be thoughtless: Social norms and individual computation. Working Papers 00-03-022, Santa Fe Institute, March 2000.
- [Griffiths and Luck, 2010] Nathan Griffiths and Michael Luck. Changing neighbours: Improving tag-based cooperation. In *Proceedings of AAMAS'10*, pages 249–256, 2010.
- [Hales and Artecconi, 2006] David Hales and Stefano Artecconi. Slacer: A self-organizing protocol for coordination in p2p networks. *IEEE Intelligent Systems*, 21:29,35, 2006.
- [Kittock, 1993] James E. Kittock. Emergent conventions and the structure of multi-agent systems. In *Proceedings of the 1993 Complex systems summer school, SFI Studies in the Sciences of Complexity Lecture Volume VI, Santa Fe Institute*, pages 507–521. Addison-Wesley, 1993.
- [Mukherjee *et al.*, 2008] Partha Mukherjee, Sandip Sen, and Stéphane Airiau. Norm emergence under constrained interactions in diverse societies. In *Proceedings of AAMAS'08 - Volume 2*, pages 779–786, Richland, SC, 2008.
- [Scott, 2000] John Scott. *Social Network Analysis: A Handbook*. Sage Publications, second. edition, 2000.
- [Sen and Airiau, 2007] Sandip Sen and Stéphane Airiau. Emergence of norms through social learning. *Proceedings of IJCAI-07*, pages 1507–1512, 2007.
- [Shoham and Tennenholtz, 1997] Yoav Shoham and Moshe Tennenholtz. On the emergence of social conventions: modeling, analysis, and simulations. *Artificial Intelligence*, 94:139–166, 1997.
- [Toivonen *et al.*, 2009] R. Toivonen, X. Castelló, V. M. Eguíluz, J. Saramäki, K. Kaski, and M. San Miguel. Broad lifetime distributions for ordering dynamics in complex networks. *Physical Review*, 79, Feb 2009.
- [Urbano *et al.*, 2009] Paulo Urbano, João Balsa, Luis Antunes, and Luis Moniz. Force versus majority: A comparison in convention emergence efficiency. pages 48–63, 2009.
- [Villatoro *et al.*, 2009] Daniel Villatoro, Sandip Sen, and Jordi Sabater-Mir. Topology and memory effect on convention emergence. In *Proceedings of the IAT*. IEEE Press, 2009.
- [Walker and Wooldridge, 1995] Adam Walker and Michael Wooldridge. Understanding the emergence of conventions in multi-agent systems. In Victor Lesser, editor, *Proceedings of AAMAS*, pages 384–389, San Francisco, CA, 1995. MIT Press.

⁸As reported on Sept 1st, 2010 at http://www.comscore.com/Press_Events/Press_Releases/2010/8/Facebook_Captures_Top_Spot_among_Social_Networking_Sites_in_India