

Achieving Coordination in Multi-Agent Systems by Stable Local Conventions under Community Networks

Shuyue Hu and Ho-fung Leung

Department of Computer Science and Engineering
The Chinese University of Hong Kong
Hong Kong, China
syhu@cse.cuhk.edu.hk and lhf@cuhk.edu.hk

Abstract

Recently, the study of social conventions has attracted much attention in the literature. We notice that a type of interesting phenomena, *local convention phenomena*, may also exist in certain multi-agent systems. When agents are partitioned into compact communities, different local conventions emerge in different communities. In this paper, we provide a definition for local conventions, and propose two metrics measuring their strength and diversity. In our experimental study, we show that agents can achieve coordination via establishing diverse stable local conventions, which indicates a practical way to solve coordination problems other than the traditional global convention emergence. Moreover, we find that with smaller community sizes, denser connections and fewer available actions, diverse local conventions emerge in shorter time.

1 Introduction

Social conventions, such as greeting rituals and wedding customs, are very common in human society. They are a type of social norms and guide people's behaviours by dictating how to act in certain situations where coordination is required. By conforming to social conventions, people's decision making processes are simplified and social conflicts between people are reduced [Pujol *et al.*, 2005].

The concept of social conventions has been widely used by the multi-agent system literature to study how coordination among agents can be achieved. To introduce social conventions into multi-agent systems, there are two branches of approaches: the prescriptive approach and the emergent approach [Savarimuthu and Cranefield, 2011]. The prescriptive one is concerned with central authorities specifying how agents should behave. The emergent one is a bottom-up approach where a social convention emerges as a natural result of agents' local interactions. Compared with the prescriptive one, it does not need a priori central authorities and is more suitable for open and dynamic systems [Delgado, 2002].

Researchers identify that the underlying network topology of multi-agent systems is a key factor of convention emergence phenomena. Prior work shows that to achieve coordination,

agents generally establish a globally accepted convention under different types of networks [Delgado, 2002; Kittock, 1993; Marchant *et al.*, 2015; Villatoro *et al.*, 2009]. A globally accepted convention, which is referred as a *global convention* in our work, is an action that is adopted by the majority (at least 90%, but not necessarily 100% of agents) voluntarily. It has now become a convention to take global conventions as the ultimate goal to pursue and mechanisms have been proposed for better achievement towards this goal [Hao and Leung, 2013; Hasan *et al.*, 2014; Shibusawa and Sugawara, 2014; Yu *et al.*, 2013].

In this paper, we study the effects of community structure on convention emergence, in which nodes are partitioned into communities. In a network with significant community structure, communities of nodes are internally densely connected, but externally loosely connected [Girvan and Newman, 2002]. Community structure widely exists in social networks, biological networks, and technological networks, and is identified as a significant property of real-life networks [Fortunato, 2010]. We hypothesize that under networks with significant community structure, the sparse inter-community connections almost isolate the agents' interactions in different communities. Thus, it should be hard for a global convention to emerge naturally from agents' interactions. However, it is unclear, a priori, if and how agents achieve coordination under such kind of networks. In human society where a number of communities exist, it is common to observe such local norm (or convention) phenomena [Cialdini and Trost, 1998]: while people of the same community reach consensus on certain behavioural standards, the standards in different communities may vary. For example, it is well recognized that drivers should drive on the same side of roads to avoid traffic accidents, however, on left or right side drivers should drive may vary in different countries. We expect that such phenomena also exist in multi-agent systems under networks with significant community structure. That is, agents of the same community establish a locally conformed social convention (*local convention*) to achieve coordination, though the local conventions may vary in different communities.

To verify our hypothesis, we define a local convention to be a restriction on behaviours, which is imposed on a community of agents. We propose two metrics, namely, conformity to local conventions in communities and diversity of agents' actions, to characterize the strength and diversity of local con-

ventions. An experimental study is conducted on community network-based multi-agent systems. Here we consider the general case where the network topology is static and all of the agents can only observe their own interactions. Moreover, it should be noticed that agents in our work are autonomous by default, as autonomy is the most important property that an agent possesses [Wooldridge and Jennings, 1995]. Our experimental results confirm the possible emergence of diverse local conventions. Such phenomena so widely exist that they almost appear in every compact community networks that we experiment on. We also investigate the individual payoff of agents in systems where diverse local conventions emerge, and identify several factors that influence the speed of local convention emergence.

The remainder of the paper is organized as follows. Section 2 reviews the related work. Section 3 introduces some basic concepts of local convention emergence. Section 4 presents our experimental study on local convention phenomena in multi-agent systems. Section 5 concludes the paper with some directions for the future work.

2 Related Work

Social conventions are regularities of behaviour which are a result of being a solution of a recurrent coordination problem, and turn normative with time [Ullmann-Margalit, 2015]. In [Shoham and Tennenholtz, 1992]’s seminal work of convention emergence, they define the notion of social conventions in a game-theoretic framework and show that it is possible for social conventions to emerge naturally from agents’ repeated interactions of pure coordination games. To evaluate the emergence of conventions, [Kittock, 1993] proposes a widely adopted criterion: a convention (or a global convention in our work) has emerged in a system if at least 90% of agents have converged to choose one same action. We refer this criterion as 90% convergence metric hereafter.

Since real-life multi-agent systems are systems in networked environments [Savarimuthu and Cranefield, 2011], a significant amount of research study how the network topology affects the convention emergence. By placing agents on regular, small world and scale free networks, [Kittock, 1993; Delgado, 2002] find that agents under networks with larger diameters spend longer time to achieve global convention emergence. [Pujol *et al.*, 2005] observe that agents can establish pareto-optimal global conventions when they are placed on random and small world networks. [Epstein, 2001; Villatoro *et al.*, 2009] identify the existence of subconventions (which refers to a small portion of agents using different actions than that adopted by the majority) under regular and scale free networks, and observe that it usually costs systems a long time to merge these subconventions into a global convention. To better dissolve the subconventions and facilitate the emergence of global conventions, [Griffiths and Luck, 2010; Villatoro *et al.*, 2011] propose to let agents to be capable of changing the network topology by rewiring their links.

In general, the prior work on effects of network topology on convention emergence [Marchant *et al.*, 2015; Mukherjee *et al.*, 2008; Sen and Sen, 2010] mostly reveals that agents under different kinds of networks can generally es-

	action a_1	action a_2
action a_1	1,1	-1,-1
action a_2	-1,-1	1,1

Table 1: Payoff Matrix of a 2-player-2-action Pure Coordination Game

ablish global conventions to achieve coordination. Local conventions, however, are considered to be usually transitory and thus to impede the convention emergence processes toward coordination [Epstein, 2001; Villatoro *et al.*, 2009; 2011]. In this work, community networks, which have not been studied in the literature before, are considered. We show that coordination can also be achieved in another feasible way, i.e. the emergence of stable diverse local conventions.

3 Local Conventions and Measures of Emergence

We follow the game-theoretic framework introduced by [Shoham and Tennenholtz, 1992], where social conventions can emerge naturally when agents are learning over the repeated pure coordination games (or precisely 2-player- m -action pure coordination games).

3.1 Pure Coordination Games

The 2-player- m -action pure coordination games are games in which there exist multiple optimal joint actions and agents have to coordinate to select among them. The two agents playing these games are rewarded if they choose the same action, and are punished if they choose different actions. In total, there are m Nash equilibria which are equally good in these games, which denotes m possible ways to achieve coordination for agents. An example of these games where $m = 2$ is shown in Table 1.

3.2 Local Conventions

Under the scenario of pure coordination games, a social convention is defined to be a “social law that restricts the agents’ behaviour to one particular strategy,” where the restriction is imposed on the majority of agents [Kittock, 1993; Shoham and Tennenholtz, 1992]. While this definition is appropriate to characterize the traditional way to achieve coordination, i.e. global convention emergence, it overlooks the possibility that agents can also achieve coordination by establishing diverse local conventions. Extending this definition, we formally define local conventions as follows.

Definition 1 (Local Convention). *Given a set N of agents, a set A of available actions for each agent, and a partition $C = \{c_1, c_2, \dots, c_m\}$ of agents such that $\forall c_i \in C, c_i \subseteq N, c_1 \cup c_2 \cup \dots \cup c_m = N, \forall c_i, c_j \in C, c_i \cap c_j = \emptyset$, a local convention λ_c of community $c \in C$ is a restriction on agents of the community c to choose a certain action $\hat{a} \in A$, denoted by $\lambda_c \rightarrow \hat{a}$.¹*

¹A local convention λ_c can also be a particular mixed strategy. Because which form of strategies agents learn is not the main interest of our work, we only focus on pure strategy here.

By Definition 1, local conventions of different communities are generally different in a multi-agent system. If they are exactly the same, then a global convention emerges indeed. Moreover, if there is only one community in a system, the local convention of this community is equivalent to the global convention of this system.

3.3 Metrics for Quantitative Measures

Conformity of Agents' Actions to Local Conventions

To measure the overall conformity to local conventions of a system of agents, we first measure the conformity to an individual local convention of a community of agents.

Intuitively, to what degree a community of agents conforms to its local convention λ_c should concern the proportion of agents using the action \hat{a} which is indicated by the convention λ_c (such that $\lambda_c \rightarrow \hat{a}$). Meanwhile, it should be irrelevant to which particular action the convention λ_c indicates, because all Nash equilibria yield the same optimal payoff in a pure coordination game. We thus propose to use Shannon entropy (or entropy hereafter) [Shannon, 2001] to measure the local convention conformity γ_c of any community $c \in C$, as follows.

Definition 2 (Local Convention Conformity γ_c of a Community c). *Given a community $c \in C$ of agents and a set A of available actions for each agent in a multi-agent system, the local convention conformity γ_c of the community c is:*

$$\gamma_c = 1 - \frac{1}{\log_2 |A|} H(a|c) \quad (1)$$

$$s.t. \quad H(a|c) = - \sum_{a \in A} p(a|c) \log_2 p(a|c)$$

where $H(a|c)$ is the entropy of agents' actions in community c , $p(a|c)$ is the relative frequency of agents choosing the action a in community c , and $\frac{1}{\log_2 |A|}$ is a normalization constant which normalizes γ_c to a scale of $[0, 1]$.

By Definition 2, the more uniform agents' actions in a community c is, the larger value of the community's local convention conformity γ_c is. If γ_c reaches the maximum, which is 1, then all of the agents inside the community c use one certain action (no matter which particular action it is). On the contrary, if γ_c reaches the minimum, which is 0, then agents inside the community c diverge to choose all of the available actions with equal relative frequencies ($\forall a \in A, p(a|c) = \frac{1}{|A|}$), indicating that no action stands out particularly.

We then aggregate the local convention conformity γ_c of each community $c \in C$ to obtain the overall conformity γ of the system to local conventions, taking the effects of community sizes into consideration, as follows:

Definition 3 (Local Convention Conformity γ of a Multi-Agent System). *Given a set C of communities in a multi-agent system, the local convention conformity (or conformity) γ of this system is:*

$$\gamma = \sum_{c \in C} p(c) \gamma_c \quad (2)$$

where $p(c)$ is the size of community c , γ_c is local convention conformity of the community c and $\gamma \in [0, 1]$.

By Definition 3, the system's overall conformity γ is the weighted sum of the conformity γ_c of each community $c \in C$. Therefore, a high degree of conformity γ (which is close to 1) indicates that almost all communities of agents fully conform to their own local conventions. By contrast, a low degree of conformity γ (which is close to 0) means that there is almost no sign of emerging local convention in each community.

Diversity in Agents' Actions

While conformity measures the overall strength of local conventions in a system, we also need to measure how diverse agents' actions throughout this system are. We again use entropy to define the diversity δ in agents' actions of a multi-agent system as follows.

Definition 4 (Diversity δ in Agents' Actions of a Multi-Agent System). *Given a set A of available actions for each agent in a multi-agent system, the diversity δ in agents' actions (diversity) of this system is:*

$$\delta = \frac{1}{\log_2 |A|} H(a) \quad (3)$$

$$s.t. \quad H(a) = - \sum_{a \in A} p(a) \log_2 p(a)$$

where $H(a)$ is the entropy of all of the agents' actions in the system, $p(a)$ is the relative frequency of agents choosing the action a , and $\frac{1}{\log_2 |A|}$ is a normalization constant which normalizes δ to a scale of $[0, 1]$.

By Definition 4, if all of the available actions are chosen by agents with equal relative frequencies ($\forall a \in A, p(a) = \frac{1}{|A|}$), which is the most diverse case of agents' actions, the diversity δ of this system will be 1. On the contrary, if all of the agents in the system choose one same action, which is the most uniform case of agents' actions, the diversity δ will be 0.

As we illustrate above, in a system with high degree in local convention conformity γ , each community of agents converge to choose one certain action. Therefore, if the diversity δ is close to 0 in such a system, it is then natural to infer that, the actions chosen by different communities are almost the same, and thus uniform local conventions (or a global convention) emerge in this system. By contrast, if the diversity δ is much larger than 0 in such a system, it is very possible that, different communities of agents converge to choose different certain actions, and thus diverse local conventions emerge in this system.

As a result, by these complementary metrics, we can have quantitative measures on the degree of global convention emergence or diverse local convention emergence in a multi-agent system.

4 An Experimental Study

We adopt social learning model [Sen and Airiau, 2007], a representative and commonly used model, for convention emergence in multi-agent systems. Under this model, agents have no prior knowledge, cannot observe others' interactions and only learn over their own repeated interactions. Specifically, in each iteration, an agent is first paired with another agent




Figure 1: An example of the effect of σ , given g, v, l, k are fixed (nodes are partitioned into 5 communities).

which is randomly selected from its neighbours. Then all of the agents choose their actions for pure coordination games independently. Later, all pairs of agents play games concurrently and receive immediate rewards or punishment. Based on what they received, agents reevaluate their choices on actions. We equip agents with Q-learning [Watkins and Dayan, 1992] with ϵ -greedy exploration. The learning rate of Q-learning is set to 0.5 and the exploration rate is set to $\epsilon = 0.1$. The number of available actions is set to 5 by default, indicating 5 different ways to achieve coordination among agents.

4.1 Underlying Network Topology

We adopt *Gaussian random partition generator* [Brandes et al., 2003] to generate the underlying network topology of multi-agent systems. This generator generates random networks with community structure (which are referred as community networks for short hereafter), such that each community is a dense Erdos-Renyi random graph, while different communities are connected randomly but sparsely.

The generation of network topology is controlled by a tuple of five parameters $\langle g, v, l, k, \sigma \rangle$. The total number of communities is given by l . The sizes of communities have a Gaussian distribution with average community size g and standard deviation v . As such, there are totally $n = g \times l$ nodes in this network. How these n nodes are connected is then controlled by the average degree k and the separation degree σ . The average degree k is the number of neighbours any node has on average. The separation degree σ , which is in the interval $(0, 1)$, specifies the relative ratio of intra-community neighbours to the total number of neighbours. That is to say, out of all the k_i neighbours of a node i , $k_i \times \sigma$ of them are expected to be in the same community of the node i . The separation degree σ characterizes the significance of community networks. When σ is small, the community structure is insignificant and the networks resemble random networks. On the contrary, when σ is large, the networks have compact communities. A graphical illustration of the effect of σ is shown in Figure 1.

In this work, we generate different types of community networks, each of which is defined by a tuple of parameters $\langle g, v, l, k, \sigma \rangle$. We select various values for these parameters so that the trends of changes can be clearly identified. Specifically, we vary the average community size g in 50, 75 and 100, the standard deviation v of community sizes in 0, 1 and 3, the number l of communities in 10, 20 and 40, the average degree k in 10, 20 and 40, and the separation degree σ from 0.1 to 0.9. In total, there are 729 types of community networks. We generate 1000 instance networks for each net-




Figure 2: Proportion of instance networks satisfying the 90% convergence metric.




Figure 3: A typical instance of the dynamics of convergence degree ($g = 50, v = 1, l = 20$ and $k = 20$).

work type. For each instance network, we conduct an one-run experiment which contains 100,000 iterations.

4.2 Initial Results on Global Conventions

Initially, we seek to find if the global conventions, which are conventionally measured by the 90% convergence [Kittcock, 1993], emerge in the instance networks. From the results presented in Figure 2, we notice a significant pattern of the convergence in agents' actions under these networks. While the 90% convergence can be generally reached in networks with small σ , the proportion of networks satisfying this metric decreases significantly when σ becomes larger. When σ equals 0.8 or 0.9, there is almost no network satisfying the metric.

We present the typical dynamics of the convergence degree (the proportion of agents choosing the most widely used action) under networks with different values of σ in Figure 3. It can be observed that while the convergence degree quickly reaches 90% when $\sigma \leq 0.7$, it is smaller than 60% and 40% when σ equals 0.8 and 0.9 respectively. Due to the lack of space, we only present our results within the first 2000 iterations. It should be noticed that the convergence degree does remain stable in all the subsequent iterations.

Our above results on the convergence in agents' actions show that agents can generally establish global conventions under networks with small σ . This can be expected, since networks with small σ resemble random networks. More importantly, our results in particular indicate that the global convention emergence is uncommon under networks with large σ (compact community networks). This interesting finding naturally raises the question that has been rarely explored before: if and how agents achieve coordination in systems where global conventions fail to emerge. To answer this question, we focus on networks with large σ ($\sigma = 0.8$ or 0.9) in the following sections.²

4.3 Emergence of Diverse Local Conventions

To confirm our hypothesis on the possible emergence of local conventions, we characterize the dynamics in instance networks using the metrics: local convention conformity γ and diversity in agents' actions δ . The statistical results are presented in Figure 4.

²It is worth mentioning that, usually, the separation degree σ is also large in real-life networks, because a great number of nodes have no links to outside their own communities [Lancichinetti et al., 2010].




Figure 4: Dynamics of conformity γ and diversity δ in instance networks with different values of σ .

We observe that when σ equals 0.8 or 0.9, conformity γ increases from around 0 to about 1 gradually in the first 500 iterations. Meanwhile, diversity δ drops from around 1 to about 0.65 when $\sigma = 0.8$, and to about 0.85 when $\sigma = 0.9$. These values of conformity and diversity then remain stable in the subsequent 99,500 iterations (the presented results limit to the first 2000 iterations due to the lack of space).

As we explain in Section 3, the high degrees in diversity and conformity complementarily indicate the emergence of diverse local conventions. That is, while each community of agents have converged to choose the same action (i.e. the community’s local convention), to which particular action different communities converge may vary. Therefore, as our results manifest, the emergence of diverse local conventions extensively exists under compact community networks, where global conventions are generally absent. Moreover, the more significant the community structure is, the more diverse the local conventions are.

Another mentionable pattern in Figure 4 is that when $\sigma \leq 0.7$, conformity γ rises to close to 1 and diversity δ drops to a low value. This result again verifies that the global convention generally emerge in networks with small σ , which has been already shown in Section 4.2.

To facilitate better understanding on local convention emergence, we present a typical example of the dynamics in systems with large σ in Figure 5. Each action is indicated by one color, the history of each agent’s actions is represented by one row and agents of different communities are partitioned for clarity. It can be observed that agents’ choices on actions seem to be very diverse in the early stage. As more iterations are experienced, agents’ actions within a community become more uniform, and a particular action pervades and dominates the community. It is interesting to note that not only the particular dominant actions of different communities are diverse, but also at what speed these actions reach dominance are different in different communities. Moreover, these dominant actions maintain their dominance in all the subsequent iterations, although there are seldom changes on agents’ actions caused by the few random explorations.

We hypothesize that the rationale of the above results can be explained as follows. Agents choose their actions randomly in the first few iterations. However, since random choices on actions cause discoordination, agents start to choose the same action as that of their opponents. The outnumbered intra-community interactions under compact com-




Figure 5: A typical instance of the dynamics in systems where global conventions fail to emerge ($g = 50$, $v = 1$, $l = 10$, $k = 20$ and $\sigma = 0.9$).

munity networks drive agents of the same community to choose the same action. The few inter-community interactions, on the other hand, almost isolate the dynamics within different communities. As a result, stable local conventions, which reveal the coordination inside communities, are born out from different communities in diverse forms and speeds.

Moreover, we also notice that when σ is large, a larger number of communities results in even more diverse local conventions. As shown in Figure 6, when σ is 0.9, a larger value in the number l of communities significantly increases the diversity, while the change of other network properties only slightly affect the diversity. The significant influence of l is reasonable. Because each community of agents establishes its own local convention under compact community networks, a larger number of communities makes the convergence of different communities more uncertain and thus increases the diversity.

Although agents independently choose their actions and only learn over their own interactions, our results confirm that agents can still extensively achieve stable coordination by establishing diverse local conventions. This is in clear contrast with prior work on convention emergence which hints that agents generally achieve coordination only via establishing a global convention, but the emergence of local conventions is generally an undesirable and transient phenomenon [Epstein, 2001; Villatoro *et al.*, 2009; 2011]. Therefore, our work in particular reveals another feasible way to achieve coordination among agents. Moreover, as our results imply, compared with the traditional global convention emergence, the local convention emergence is a more practical way to achieve coordination for agents under compact community networks (with large σ).

4.4 Payoffs in Diverse Local Convention Emergence

We now evaluate how agents’ payoffs are affected by the emergence of diverse local conventions. We present the dynamics of agents’ average individual payoff under networks with large σ in Figure 7. It can be seen that when σ equals 0.8 or 0.9, the payoff is approximately -0.5 in the first few iterations. Then as the number of iterations increases, the payoff grows gradually to about 0.5 when $\sigma = 0.8$, and to around 0.56 when $\sigma = 0.9$. Interestingly, by comparing Figure 7 and Figure 4, we observe that the rising trend in the average




Figure 6: Effects of parameter change on diversity δ (baseline: $g = 50, v = 1, l = 10, k = 20$ and $\sigma = 0.9$).




Figure 7: Dynamics of average individual payoff under networks with large separation degree σ .

individual payoff is consistent with the increase of local convention conformity.


We explain the dynamics of agents' payoffs as follows. In the first few iterations, because of agents' random choices on actions, most agents are in discoordination and receive much punishment. When local conventions start to emerge in communities, agents are increasingly in coordination with their intra-community neighbours, which yields a large increase in their payoffs. However, since the local conventions are generally diverse, agents may be still in discoordination with their few inter-community neighbours. Thus, under a more compact community network which features fewer inter-community neighbours, agents' payoffs rise with the gradual emergence of local conventions and stabilize at a higher level.

4.5 Speed of Diverse Local Convention Emergence


We investigate the speed of local convention emergence under compact community networks by measuring at which iteration conformity γ begins to become stable.

Let us first consider the effects of network properties. Obviously, the average community size g and average degree k are network properties that significantly influence the size and the interior structure of individual communities. The standard deviation v of community sizes and the number l of communities, on the other hand, are irrelevant to the individual communities. Interestingly, from our results presented in Figures 8(a)-8(d), we identify a significant difference in the influence of the above two classes of network properties. With smaller average community size or higher average degree, agents can achieve local coordination within fewer iterations. In contrast, the standard deviation of community sizes and the number of communities have little influence on the speed of local convention emergence.


We conjecture that such difference is again caused by the few inter-community interactions under compact community networks. Because the few inter-community interactions almost isolate the emergence processes in different communities, how fast local conventions emerge mainly concerns the networks properties which affect individual communities. A larger average size of communities introduces more uncertainty into the overall agents' choices on actions and thus slows down the local convention emergence processes. A higher average degree, which leads to more interactions in each iteration, facilitates agents to achieve coordination us-



(a) average community size g (b) standard deviation of community sizes v



(c) the number of communities l (d) average degree k



(e) the number of available actions m

Figure 8: Effects of network properties and the number of available actions on the speed of diverse local convention emergence.

ing fewer iterations.

Next, we consider the effect of the number of available actions. Intuitively, the increase in the number of available actions diverges an agent's choice on actions. Therefore, as it is confirmed in Figure 8(e), agents with more available actions spend more iterations in establishing local conventions.

One of the implications from our results is that, the time required for achieving coordination under compact community networks mainly depends on the community sizes, the link density and the number of available actions.

5 Conclusion and Future Work

In this paper, we define a local convention to be a restriction of behaviour, which is imposed on a community of agents. We provide two quantitative measures: local convention conformity and diversity in agents' actions, to measure the strength and diversity of local conventions. The main result of our experimental study is that the stable coordination among agents is achieved via diverse local convention emergence under compact community networks, where the traditional global convention emergence is almost absent. Also, as our results indicate, with the gradual emergence of local conventions, agents payoffs increase. Moreover, smaller community sizes, denser connections between agents and fewer available actions can lead to faster emergence of local conventions. In our future work, we will investigate how overlapping community structure, more unbalanced community size and degree distribution influence the local convention emergence. In addition, we will consider the cases where fixed strategy agents are introduced, agents have the ability to rewire their links, and agents are aware of with whom they interact, to see if diverse stable local conventions still emerge.

References

- [Brandes *et al.*, 2003] Ulrik Brandes, Marco Gaertler, and Dorothea Wagner. Experiments on graph clustering algorithms. In *European Symposium on Algorithms*, pages 568–579. Springer, 2003.
- [Cialdini and Trost, 1998] Robert B Cialdini and Melanie R Trost. *Social influence: Social norms, conformity and compliance*. McGraw-Hill, 1998.
- [Delgado, 2002] Jordi Delgado. Emergence of social conventions in complex networks. *Artificial Intelligence*, 141(1):171–185, 2002.
- [Epstein, 2001] Joshua M Epstein. Learning to be thoughtless: Social norms and individual computation. *Computational Economics*, 18(1):9–24, 2001.
- [Fortunato, 2010] Santo Fortunato. Community detection in graphs. *Physics Reports*, 486(3):75–174, 2010.
- [Girvan and Newman, 2002] Michelle Girvan and Mark EJ Newman. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences*, 99(12):7821–7826, 2002.
- [Griffiths and Luck, 2010] Nathan Griffiths and Michael Luck. Norm emergence in tag-based cooperation. page 80, 2010.
- [Hao and Leung, 2013] Jianye Hao and Ho-fung Leung. The dynamics of reinforcement social learning in cooperative multiagent systems. In *Proceedings of the 23rd International Joint Conferences on Artificial Intelligence*, volume 13, pages 184–190, 2013.
- [Hasan *et al.*, 2014] Mohammad R Hasan, Sherief Abdallah, and Anita Raja. Topology aware convention emergence. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 1593–1594. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
- [Kittock, 1993] James E Kittock. Emergent conventions and the structure of multi-agent systems. In *1993 lectures in Complex Systems: the Proceedings of the Complex Systems Summer School*, 1993.
- [Lancichinetti *et al.*, 2010] Andrea Lancichinetti, Mikko Kivela, Jari Saramaki, and Santo Fortunato. Characterizing the community structure of complex networks. *PLoS one*, 5(8):e11976, 2010.
- [Marchant *et al.*, 2015] James Marchant, Nathan Griffiths, and Matthew Leke. Convention emergence and influence in dynamic topologies. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1785–1786, 2015.
- [Mukherjee *et al.*, 2008] Partha Mukherjee, Sandip Sen, and Stéphane Airiau. Norm emergence under constrained interactions in diverse societies. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent systems-Volume 2*, pages 779–786, 2008.
- [Pujol *et al.*, 2005] Josep M Pujol, Jordi Delgado, Ramon Sangüesa, and Andreas Flache. The role of clustering on the emergence of efficient social conventions. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, pages 965–970, 2005.
- [Savarimuthu and Cranefield, 2011] Bastin Tony Roy Savarimuthu and Stephen Cranefield. Norm creation, spreading and emergence: A survey of simulation models of norms in multi-agent systems. *Multiagent and Grid Systems*, 7(1):21–54, 2011.
- [Sen and Airiau, 2007] Sandip Sen and Stéphane Airiau. Emergence of norms through social learning. In *Proceedings of 20th International Joint Conferences on Artificial Intelligence*, volume 1507, page 1512, 2007.
- [Sen and Sen, 2010] Onkur Sen and Sandip Sen. Effects of social network topology and options on norm emergence. In *Coordination, Organizations, Institutions and Norms in Agent Systems*, pages 211–222. Springer, 2010.
- [Shannon, 2001] Claude Elwood Shannon. A mathematical theory of communication. *ACM SIGMOBILE Mobile Computing and Communications Review*, 5(1):3–55, 2001.
- [Shibusawa and Sugawara, 2014] Ryosuke Shibusawa and Toshiharu Sugawara. Norm emergence via influential weight propagation in complex networks. In *Network Intelligence Conference (ENIC), 2014 European*, pages 30–37. IEEE, 2014.
- [Shoham and Tennenholtz, 1992] Yoav Shoham and Moshe Tennenholtz. Emergent conventions in multi-agent systems: Initial experimental results and observations. *KR-92*, pages 225–231, 1992.
- [Ullmann-Margalit, 2015] Edna Ullmann-Margalit. *The emergence of norms*, volume 11. Oxford University Press, USA, 2015.
- [Villatoro *et al.*, 2009] Daniel Villatoro, Sandip Sen, and Jordi Sabater-Mir. Topology and memory effect on convention emergence. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology-Volume 02*, pages 233–240, 2009.
- [Villatoro *et al.*, 2011] Daniel Villatoro, Jordi Sabater-Mir, and Sandip Sen. Social instruments for robust convention emergence. In *Proceedings of the 22nd International Joint Conferences on Artificial Intelligence*, volume 11, pages 420–425, 2011.
- [Watkins and Dayan, 1992] Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine Learning*, 8(3-4):279–292, 1992.
- [Wooldridge and Jennings, 1995] Michael Wooldridge and Nicholas R Jennings. Intelligent agents: Theory and practice. *The knowledge engineering review*, 10(02):115–152, 1995.
- [Yu *et al.*, 2013] Chao Yu, Minjie Zhang, Fenghui Ren, and Xudong Luo. Emergence of social norms through collective learning in networked agent societies. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multiagent Systems*, pages 475–482, 2013.