The Effects of Fixed-Strategy Agents on Local Convention Emergence in Multi-agent Systems

Tim Borglund, Shuyue Hu, Ho-Fung Leung

► To cite this version:

Tim Borglund, Shuyue Hu, Ho-Fung Leung. The Effects of Fixed-Strategy Agents on Local Convention Emergence in Multi-agent Systems. 10th International Conference on Intelligent Information Processing (IIP), Oct 2018, Nanning, China. pp.99-108, 10.1007/978-3-030-00828-4_11 . hal-02197792

HAL Id: hal-02197792

<https://inria.hal.science/hal-02197792>

Submitted on 30 Jul 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed under a Creative Commons Attribution| 4.0 International License

The Effects of Fixed-Strategy Agents on Local Convention Emergence in Multi-Agent Systems

Tim Borglund¹, Shuyue Hu² and Ho-Fung Leung²

¹ Department of Computer Science and Engineering, Lund University, Lund, Sweden
dic13tbo@student.lu.se

² Department of Computer Science and Engineering, The Chinese University of Hong Kong,
Hong Kong, China
{syhu, lhf}@cse.cuhk.edu.hk

Abstract. Achieving coordination in multi-agent systems has previously been found to be possible by utilizing local conventions as opposed to relying on the emergence of global conventions. On another note, fixed-strategy agents have been researched to manipulate the behaviour of networks with global conventions, but not local conventions. This paper studies how fixed-strategy agents impact local convention emergence and if they could be useful for both compact and loose community structures. It is shown that while the existence of a larger number of fixed-strategy agents generally makes local conventions emerge faster, only a few fixed-strategy agents are needed to convince communities to use their fixed action. Finally, fixed-strategy agents are helpful for compact community networks but not for loose community networks.

Keywords: Multi-agent systems, intelligent agents, local convention emergence, fixed-strategy agents.

1 Introduction

Conforming to social conventions is an important factor when people make decisions. People want to fit into their communities or countries by following unwritten rules. These could for example be to wait in line or using the correct greeting gesture. In order for people within a community to live together effectively, following conventions is socially required.

When it comes to artificial intelligence, figuring out how conventions can be followed by agents in multiagent systems has been researched in many different ways. Often this relates to finding a global convention that most of agents in a network have to follow in order to achieve coordination. For an agent to follow a global convention it means that it performs the same action as most of the other agents in the network. Hu and Leung (2017) propose another way to achieve coordination. The network is split into communities and the goal for the agents is to follow so-called local conventions. This means that an agent needs to follow the same convention as the other agents within the same community, but not necessarily with agents from other communities. Different cultures and countries in the world often have different conven-

tions. Therefore, finding local conventions in a network could be used to give a more accurate representation of these differences, as opposed to finding global conventions.

Another topic that has been researched is how fixed-strategy agents can be used to influence the normal (learning) agents of a network (Griffiths and Anand 2012; Marchant, Griffiths and Leeke 2014). Fixed-strategy agents are agents that always perform the same action, instead of learning what action to perform in order to follow a convention. In a sense, they can be seen as being stubborn teachers who try to teach the other agents what convention to follow, without learning anything themselves.

The work of this paper is primarily based on the work that aforementioned Hu and Leung (2017) have done with regards to local conventions. Continuing from there, this paper shows how fixed-strategy agents can be used to impact local convention emergence under community structures. This raises a number of interesting questions. How many fixed-strategy agents are needed to convince a community to conform to a certain convention, how they can affect the speed at which local conventions emerge, whether the placement of the fixed-strategy agents makes any difference to this speed, and if they could be useful for both compact and loose community structures. These questions will be investigated in this paper.

This paper is structured as follows: Section 2 describes related work that has been done in relation to fixed-strategy agents and local convention emergence in multi-agent systems. Section 3 describes which underlying methods and definitions have been used to conduct the experiments. Section 4 shows how the experiments are performed and what the results are. Section 5 gives the conclusions of the paper and future work is suggested.

2 Related Work

How fixed-strategy agents impact and manipulate multi-agent systems has been researched lately. Initially this was experimented on by Sen and Airiau (2007) when proposing the social learning model. They noticed that only a small number of fixed-strategy agents were needed in order to convince the other agents of the network to use their fixed action as a global convention. Griffiths and Anand (2012) showed how fixed-strategy agents affect the speed at which global conventions emerge and also how this was affected by inserting them with different placement strategies. Marchant, Griffiths and Leeke (2014) showed how fixed-strategy agents can be used to destabilize already existing global conventions.

Communities within a network are clusters of nodes that are internally compactly connected, but loosely connected to other nodes of the network (Girvan and Newman, 2002). In real-life societies where communities exist, it has been shown that local conventions are common (Cialdini and Trost, 1998).

When it comes to achieving coordination in multi-agent systems, most of the work has been focused on methods related to reaching a global convention. However, as stated in the introduction, Hu and Leung (2017) found that achieving coordination in multi-agent systems does not necessarily have to be done through global conventions. It can also be done by finding local conventions within communities of the network.

This offers an alternative way of achieving coordination that is more flexible since all of the agents in the network do not have to perform the same action. They defined compact community networks where the agents mostly have connections with agents of the same community and found that local conventions can arise in these networks.

Consequently, the notion of fixed-strategy agents and local conventions in community structures are separately nothing new. However, the combination of these has not been studied before and is the topic of this paper.

3 System Model

The system model used for the experiments in this paper is in general similar to the ones used by Hu and Leung (2017) in their paper. The difference is the inclusion of a new parameter which indicates the ratio of fixed-strategy agents there are in the network, how the fixed-strategy agents are placed in the network, and how the experiments are run. The subsections of this section will explain how their methods are used in combination with the fixed-strategy agents.

3.1 Social learning model and the pure coordination game

The learning model that has been used for the agents is the social learning model proposed by Sen and Airiau (2007). Using this model, agents learn by interacting with one of their randomly selected neighbours, which is done repeatedly a certain number of iterations. With its randomly selected neighbour, the agent will play a pure coordination game, which was introduced by Shoham and Tennenholtz (1992). The game is a 2-player-m-action coordination game where the agents are rewarded if they choose the same action and punished if they choose different actions. The payoff is the same for each of the actions, resulting in m different Nash equilibria.

The agents have no prior knowledge and cannot observe what action other agents of the network are choosing. After each iteration, agents update their values for the different actions. Updating the values will be done using Q-Learning (Watkins and Dayan 1992) with ϵ -greedy exploration. The fixed-strategy agents do not learn anything from the game, they simply always use their fixed action.

3.2 Local conventions

The difference between global conventions and local conventions is that with local conventions the agents in the network can achieve coordination without all of them necessarily using the same action. This is accomplished by splitting the network into communities and having a local convention for each community. The definition for local conventions is given by Hu and Leung (2017) and is structured as follows:

Given a set of N agents, a set of A 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 γ_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 $\gamma_c \rightarrow \hat{a}$.

3.3 Conformity

In order to know when a network has achieved coordination, a measure called conformity γ has been used. Conformity γ , as specified by Hu and Leung (2017), is a measurement of how well the agents conform to the most dominant action in their respective communities. Conformity γ is on a scale of 0 to 100%. If all of the agents perform the same action as the other agents in their respective communities, γ will be 100% and if all of them perform different actions than the other agents in their respective communities, γ will be 0%. Both the normal agents and the fixed-strategy agents are included in the conformity. With the measurement γ , we set the criterion of achieving coordination in a network to be $\gamma \geq 90\%$.

3.4 Network topology

Generating the networks that are used for the experiments is done using a Gaussian random partition generator (Brandes, Gaertler and Wagner 2003), with the addition of a new parameter for the fixed-strategy agents as well as how to place them. The generator generates networks that have community structure, meaning that the agents are divided into different clusters (communities). Each community will be an Erdős-Rényi random graph and generally have many connections between the agents within the community, and not so many to other agents of the network (depending on the separation degree which is explained in the list below).

There are a number of different parameters that goes into the generator:

- a . The number of available actions for each agent to choose from.
- g . The average size of each community.
- v . The standard deviation of the size of each community.
- l . The number of communities.
- σ . The separation degree. This is a value between 0 and 1 which indicates how loose or compact the communities should be. A separation degree of 0 means that there is no community structure and a separation degree of 1 means that agents only have connections with agents of their own community. Higher values therefore represent compact community networks and lower values represent loose community networks.
- k . The average number of neighbours for an agent. The probability of being a neighbour with one of the agents in the same community is $p_{in} = \frac{k\sigma}{g-1}$ and with one of the agents in another community is $p_{out} = \frac{k-k\sigma}{n-g}$.
- The percentage of fixed-strategy agents. This is a value between 0 and 1 and shows the ratio of fixed-strategy agents in the network compared to normal agents. A value of 0 means that there are zero fixed-strategy agents and a value of 1 means that every agent is a fixed-strategy agent.

3.5 Placement of fixed-strategy agents

Placing the fixed-strategy agents is done by first generating the network topology and then replacing some of the normal agents with fixed-strategy agents. Each community will be given an equal number of fixed-strategy agents and the fixed-strategy agents within a community will perform the same fixed action. The fixed action is random for each community. The fixed-strategy agents are going to be placed according to three different placement strategies:

- Random. This strategy simply randomly replaces normal agents with fixed-strategy agents within each community.
- Most neighbours. This strategy replaces the normal agents that have the most number of neighbours within each community with fixed-strategy agents.
- Fewest neighbours. This strategy replaces the agents that have the fewest number of neighbours within each community with fixed-strategy agents.

4 Experimental Results

This section shows the results that have been found using the fixed-strategy agents in different aspects. The experiments are conducted using a combination of the parameters given in Section 3.3 by running the simulations 1,000 times and taking the average value. If nothing else is stated, the default values for the parameters are $a = 5$, $g = 100$, $v = 1$, $l = 10$, $k = 10$ and $\sigma = 0.9$. The learning rate of the Q-learning is set to 0.5 and the exploration rate is set to 0.1. The default placement strategy is random placement.

4.1 The speed of local convention emergence

One of the interesting questions regarding fixed-strategy agents in networks with community structure is how effective they can be to increase the speed at which local conventions emerge, i.e. when the conformity γ of the network reaches 90%. This is illustrated in Fig. 1. As can be seen, the general trend is that more fixed-strategy agents reduce the amount of iterations needed until local conventions emerge.

What is interesting to note in the figure is that the fixed-strategy agents do not always help the local conventions emerge faster. When the population includes only a few fixed-strategy agents, in this case 1 – 2%, the process is actually slowed down. It is not obvious why this would happen. In order to understand this, the experiments of the next section are run to see if the fixed-strategy agents manage to convince the communities to use their fixed action.

Fig. 1. The figure shows how the speed at which local conventions emerge changes when the percentage of fixed-strategy agents is varied.

4.2 Convincingness of the fixed-strategy agents

Another interesting question is if, and in that case when, the fixed-strategy agents are convincing enough for the communities to use their fixed action. This can be verified by seeing if the local convention that each community ended up conforming to is the same as the action that the fixed-strategy agents in that community are using. This is illustrated in Fig. 2. The average percentage of communities that are convinced increases as the number of fixed-strategy agents increase. With this particular set of parameters, already when the population includes 4% of fixed-strategy agents, 90% of the communities are convinced. With 7%, all of the communities are convinced. Exactly at which percentages this happens varies depending on the other parameters, but the general trend is the same.

Note that at 0% of fixed-strategy agents, the figure shows that 20% of the communities are convinced. This might seem counterintuitive, but the reason is that the number of available actions for this simulation is 5 and therefore one fifth, i.e. 20%, of the communities would have randomly chosen to use the same action as the fixed strategy agents would have been using. The curve therefore starts at 20% to give a more accurate representation of the usefulness of the fixed-strategy agents throughout the graph.

With this knowledge, the uncertainty from Section 4.1 can be answered. The reason why low percentage of fixed-strategy agents can slow down the process of finding local conventions is that they are too few to convince the communities to use their fixed action. Instead, they teach agents to use an action that will later not be used as the local convention. Thus they ‘spread misinformation’ and slow down the process of normal agents’ learning of which action should be used.

Fig. 2. The figure shows the percentage of communities that ended up using the same action as the fixed-strategy agents in that community, while varying the percentage of fixed-strategy agents.

4.3 Placement Strategies

As described in Section 3.3, the fixed-strategy agents will be placed according to three different strategies. The result is shown in Fig. 3. As can be seen, the different placement strategies do not have much of an impact on the speed of local convention emergence. Given how k , the average number of neighbours for each agent, is calculated for these kind of random networks there is not a large difference between the number of neighbours that the agents have. For the default network specified by this paper, where $k = 10$, the agent with the least number of neighbours usually has 2-4 neighbours, and the agent with the most number of neighbours usually has 16-18 neighbours, with most of the agents having 8-12 neighbours.

It therefore becomes clear that the variance in the number of neighbours is too small to see a significant difference between the strategies. The placement strategies could have a greater impact if other kinds of networks were used instead. For example for scale free networks, where a few nodes have a large number of neighbours and the rest have hardly any neighbours, selecting the few agents with a large number of neighbours as fixed-strategy agents would result in a more powerful impact on the network.

Fig. 3. The figure shows the speed of local convention emergence using different placement strategies.

4.4 Varying the separation degree

As stated in Section 3.3, the separation degree states how loose or compact the communities should be. Fig. 4 shows what percentage of instances (simulations) that manages to reach conformity γ of 90% within 100,000 iterations while varying the percentage of fixed-strategy agents and having different values for the separation degree. For high values of the separation degree, i.e. 0.8 and 0.9, all of the instances reaches 90% conformity regardless of how many fixed-strategy agents there are. The fixed-strategy agents are in these cases therefore useful for manipulating the local conventions.

Fig. 4. The figure shows the percentage of instances (simulations) that ended up reaching 90% conformity after 100,000 iterations while varying the number of fixed-strategy agents and the separation degree.

However, for separation degrees of 0.6 and 0.5, the curve drops steeply and at around 10% of fixed-strategy agents almost zero of the instances reached a conformi-

ty of 90%. It should be noted that similar curves are shown for even lower separation degrees as well. Intuitively this makes sense since looser community structures means that agents are going to communicate a lot with agents from other communities. This means that they are going to be affected by teachers from different communities with different fixed actions. This will make it difficult for any local conventions to emerge. It is possible that local conventions emerge eventually, but more than 100,000 iterations can be considered an unreasonable amount.

An interesting case is when the separation degree is 0.7. The curve follows the lower separation degrees in the beginning but at around 15% of fixed-strategy agents the curve starts to go upwards again. Exactly why this happens is not clear. We hypothesize that the reason is that with enough fixed-strategy agents and still an enough level of compactness of the communities, the fixed-strategy agents are just convincing enough to the agents of their own community while not interfering too much with agents of the other communities in order for local conventions to emerge in the communities. Further research would be needed to explain this with more certainty.

5 Conclusions and Future Work

The work performed in this paper has investigated how fixed-strategy agents affect the local convention emergence in networks with community structures. One of the conclusions is that more fixed-strategy agents will make the local conventions emerge faster. This is however not always true. For very small amounts of fixed-strategy agents (around 1% to 3%) the local conventions will emerge slower. Also, only a few fixed-strategy agents (around 4%) are enough to convince almost all of the communities to conform to their fixed action. At about 7%, all of the communities are convinced. Exactly at which percentages these events happen depends on the parameters of the network, but the general tendency is the same.

Placing fixed-strategy agents randomly or based on the agents that have the most or least number of neighbours does not make any significant difference to the speed at which local conventions emerge. This is valid for random graphs generated by the Gaussian random partition generator when the number of neighbours that each agent has do not vary too much.

For compact community networks (with separation degrees of 0.8 or 0.9) the fixed-strategy agents are helpful and make the communities conform faster. For networks with loose community structure (with separation degrees of 0.6 or less) the fixed-strategy agents will make it impossible for local conventions to emerge in the communities.

There are multiple ways in which future work can be conducted. Increasing the variance of the number of neighbours that each node has for the networks used in this paper, or using some other type of network, e.g., scale free network, is one. Another continuation could be to explore exactly what happens with the

network when it is on the edge of having compact or loose community networks, i.e. when the separation degree is 0.7. Yet another interesting continuation could be to see what happens when fixed-strategy agents are introduced after some iterations have already gone by, or removing them after some iterations.

References

1. Barabási A., Albert R., “Emergence of scaling in random networks”, *Science*, vol. 286 issue 5439, 1999, p. 509–512.
2. Brandes U., Gaertler M., Wagner D., “Experiments on graph clustering algorithms”, *European Symposium on Algorithms ESA 2003*, p. 568-579.
3. Griffiths N., Anand S., “The Impact of Social Placement of Non-Learning Agents on Convention Emergence”, *11th International Conference on Autonomous Agents and Multiagent Systems AAMAS-2012*, p. 1367-1368.
4. Kittock J., “Emergent conventions and the structure of multi-agent systems”, *1993 Complex systems summer school, Santa Fe Institute Studies in the Sciences of Complexity Lecture*, p. 507-521.
5. Marchant J., Griffiths N., Leeke M., “Destabilising Conventions: Characterising the Cost”, *8th International Conference on Self-Adaptive and Self-Organizing Systems SASO-2014*, p. 139-144.
6. Girvan M., Newman M., “Community structure in social and biological networks”, *Proceedings of the National Academy of Sciences*, 2002, 99(12):7821–7826.
7. Cialdini R., Trost M., “Social influence: Social norms, conformity and compliance”, 1998, McGraw-Hill.
8. Sen S., Airiau S., “Emergence of norms through social learning”, *20th International Joint Conferences on Artificial Intelligence IJCAI-07*, p. 1507-1512.
9. Shoham Y., Tennenholtz M., “Emergent conventions in multi-agent systems: Initial experimental results and observations”, *3rd International Conference on Principles of Knowledge Representation and Reasoning KR '92*, p. 225-231.
10. Watkins C., Dayan P., “Q-learning”, *Machine Learning*, vol. 8 no. 3-4, 1992, p. 279-292.