

Convention Emergence and Influence in Dynamic Topologies

(Extended Abstract)

James Marchant, Nathan Griffiths and Matthew Leeke
Department of Computer Science
University of Warwick
Coventry, UK
{james, nathan, matt}@dcs.warwick.ac.uk

ABSTRACT

Coordination is essential to the effective operation of multi-agent systems. Convention emergence offers a low-cost and decentralised method of ensuring compatible actions and behaviour, without requiring the imposition of rules dictating such. In this paper, we explore the emergence of conventions within a dynamic network, and examine the effectiveness of fixed strategy agents in this context. We consider the importance of the dynamic nature of the system and introduce a novel heuristic, LIFE-DEGREE, to explore this.

Categories and Subject Descriptors

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

General Terms

Experimentation, Algorithms

Keywords

Dynamic networks; Conventions; Social Norms; Influence

1. INTRODUCTION

Within multi-agent systems cooperation and coordination of actions and goals enables efficient interactions. Incompatible actions result in clashes that often incur a resource cost, such as time, to the participating agents. It is often difficult to pre-determine which actions clash and so the emergence of conventions, in the form of socially-adopted expected behaviours, are utilised.

Conventions have been shown to emerge naturally from local agent interactions and enhance agent coordination by placing *social constraints* on agents' action choices. They emerge in systems with minimal assumptions, namely agent rationality and access to a (limited) memory of previous interactions to enable agents to learn.

Sen and Airiau [5] investigated social learning for convention emergence, where agents receive a payoff from their interactions and use this to inform their learning. They

showed that convention emergence can occur when agents have no memory of interactions and are only able to observe their own rewards. However, their model is limited in that agents are able to interact with any other member of the population. In many domains agents will be situated in a network topology that limits their interactions to their neighbours. Griffiths and Anand [2] show that *where* the fixed strategy agents are placed within this topology increases their effectiveness. Placement by heuristics such as degree gives improved convention emergence over random placement.

Often the nature of the relationships between agents is not static. Agents may leave or join the system, and the links between agents may change over time. These dynamic interaction topologies induce different system characteristics than those found in static networks. We consider convention emergence within these topologies and their manipulation via fixed strategy agents. We explore the effect of agent age on the effectiveness of fixed strategy agent selection and introduce a new heuristic, LIFE-DEGREE, to enable this.

2. CONVENTION EMERGENCE MODEL

We utilise a particle-based simulation, developed by González et al. [1], to model dynamic networks with preferential attachment. Agents are represented as colliding particles and the topology is modified by collisions creating an edge between agents. An agent's speed is proportional to their degree. Higher degree nodes thus have an increased probability of further collisions, further increasing their degree.

Each timestep, every agent chooses one of its neighbours in the network at random. These agents play a round of the n-action pure coordination game choosing the action they believe to be best. Agents do not know what their opponent has chosen. The payoff that each agent receives depends on the combination of chosen actions: if both chose the same action, they receive a positive payoff (+4); if the actions differ they receive a negative payoff (-1). Each agent monitors their expected payoff for each action and updates their associated utility using a simplified form of Q-Learning. A convention is considered to have emerged when 90% of non-fixed strategy agents would choose it.

We consider two different scenarios: placing fixed strategy agents at the beginning of a system's life, to encourage and direct initial convention emergence in a population, and inserting fixed strategy agents once a convention has emerged to attempt to change it. In the former case, the fixed strat-

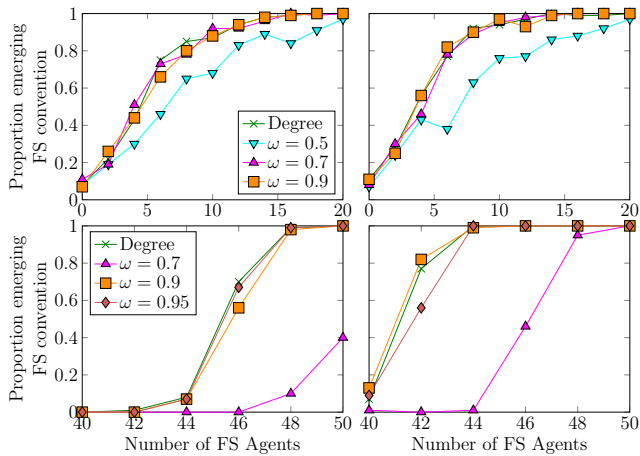


Figure 1: The proportion of runs where the fixed strategy action emerges as convention. The top row contains initial interventions, the bottom contains late interventions. The left column contains static heuristics, the right contains updating heuristics.

egy will be randomly chosen from the available actions. In the latter, it will be randomly chosen from the available actions excluding the already established convention.

To examine the effect of dynamic network characteristics on convention emergence we introduce a new placement heuristic, LIFE-DEGREE, that allows us to consider agent age and examine its importance. LIFE-DEGREE is defined as: $\omega \times deg_{norm}(n) + (1 - \omega) \times E_{rTTL}(n)$ where $deg_{norm}(n)$ is the degree of a node normalised against the largest connected component, $E_{rTTL}(n)$ is the expected remaining life of the node and ω is a weight. We also consider two forms of pure degree-based placement heuristic: Static Degree, where the agents are chosen only once (except for replacement of expired agents) and Updating Degree where the agents are updated each timestep. LIFE-DEGREE is also used in static and updating forms.

3. RESULTS

We first consider whether conventions emerge in dynamic topologies with no interventions. The size of the action space (10) determines the proportion of times a strategy emerges as a convention (≈ 0.1). With initial intervention we find that, introducing only a few fixed strategy agents at targeted locations, enables manipulation of the emerged convention more than 50% of the time. These results are shown in the first row of Figure 1. Importantly, there is very little difference between the two heuristics. Updating Degree (in the right column) slightly outperforms Static Degree (in the left) but, given the additional complexity and resource requirements needed for the updating heuristic, Static Degree would likely be sufficient in most cases.

Having established that convention emergence and its manipulation are possible, we examine the effect of considering agent age, using LIFE-DEGREE. We find that, when given equal weighting between expected life and degree ($\omega = 0.5$), LIFE-DEGREE performs markedly worse than the equivalent pure degree heuristic. This is because such a weighting ends up being heavily biased to much younger agents. A weighting of 0.7 exhibits similar performance to Static De-

gree and further increasing the weighting gave no additional improvement. These results show that an agent’s connectivity (i.e. degree) is a larger component of its influence than the amount of time it will remain in the system. The fact that considering age can only decrease the effectiveness of the chosen agents indicates that agents’ short-term influence is a larger factor in convention emergence than choosing long-term targets.

We now look to the related use of fixed strategy agents in *destabilising* and replacing an already established convention [3, 4]. This requires a convention to have emerged within the system which we allow to naturally occur without the use of fixed strategy agents. It was found that conventions always emerged before timestep $t = 1500$ and, as such, insertion of fixed strategy agents occurs at this time. This also means that the system will have entered the quasi-stationary state [1], and the topology and convention can be considered truly emerged. The results for this are shown in the bottom row of Figure 1.

Similarly to initial intervention consideration of age is detrimental to fixed strategy agent selection. In fact, the value of ω required before results become comparable to pure degree is much higher. This indicates that destabilisation is even more sensitive to the consideration of agent longevity than initial convention emergence. Considering agents of high degree is most effective at influencing convention emergence, regardless of how long that agent lasts. This is reinforced by the difference between the static and updating heuristics; such a difference was not present in initial intervention. The importance of choosing a high degree agent is far more pronounced and benefits from up-to-date information as indicated by the superior performance of the updating heuristics.

4. CONCLUSIONS

We have shown that convention emergence is possible in dynamic networks and that, similarly to static networks, it can be influenced by using fixed strategy agents. We examined the efficacy of degree-based placement heuristics when applied in both static and updating manners and found they offered improvements over random placement. We also explored the effect of considering an agent’s age and showed that consideration purely of degree was beneficial. Finally we explored these aspects when applied to convention destabilisation. We showed there were many parallels between initial and late intervention but that late intervention was more sensitive to up-to-date information within the system.

REFERENCES

- [1] M. C. González, P. G. Lind, and H. J. Herrmann. System of mobile agents to model social networks. *Phys. Rev. Lett.*, 96:088702, 2006.
- [2] N. Griffiths and S. S. Anand. The impact of social placement of non-learning agents on convention emergence. In AAMAS’12, pages 1367–1368, 2012.
- [3] J. Marchant, N. Griffiths, and M. Leeke. Destabilising conventions: Characterising the cost. In SASO’14, 2014.
- [4] J. Marchant, N. Griffiths, M. Leeke, and H. Franks. Destabilising conventions using temporary interventions. In COIN’14, 2014.
- [5] S. Sen and S. Airiau. Emergence of norms through social learning. In IJCAI’07, pages 1507–1512, 2007.