

Limited Observations and Local Information in Convention Emergence

(Extended Abstract)

James Marchant
University of Warwick
Coventry, CV4 7AL, UK
james@dcs.warwick.ac.uk

Nathan Griffiths
University of Warwick
Coventry, CV4 7AL, UK
nathan@dcs.warwick.ac.uk

ABSTRACT

In multi-agent systems it is often desirable for agents to adhere to standards of behaviour that minimise clashes and wasting of (limited) resources. In situations where it is not possible or desirable to dictate these standards globally or via centralised control, convention emergence offers a lightweight and rapid alternative. Placing fixed strategy agents within a population has been shown to facilitate faster convention emergence with some degree of control. Placing these fixed strategy agents at topologically influential locations (such as high-degree nodes) increases their effectiveness. However, finding such influential locations often assumes that the whole network is visible or that it is feasible to inspect the whole network in a computationally practical time, a fact not guaranteed in many real-world scenarios. We present an algorithm, PO-PLACE, that finds influential nodes given a finite number of network observations. We show that PO-PLACE finds sets of nodes with similar reach and influence to the set of high-degree nodes and we then compare the performance of PO-PLACE to degree placement for convention emergence in several real-world topologies.

Keywords

Convention Emergence; Partial Observability; Local Information

1. INTRODUCTION

Choosing coordinated actions allows independent agents in multi-agent systems (MAS) to increase efficiency and avoid incompatible action clashes that may incur resource waste. Establishing interaction protocols and behaviour helps avoid these clashes. However, it is often undesirable or impossible to do this in a centralised way, particularly in open MAS where agents may be owned by multiple parties or the system complexity makes it infeasible. Additionally, designers may not wish to make the choice for agents if no evidently optimal action selection exists.

Convention emergence allows a system to deal with these issues by a preference amongst all agents emerging organically in a decentralised, online manner. A convention rep-

Table 1: Summary of PO-Place performance. The sum of degrees of the selected locations is given as a proportion of that of pure degree placement.

	$o = 500$		$o = 3500/5000$	
	$n = 5$	$n = 30$	$n = 5$	$n = 30$
CondMat	0.90 (0.05)	0.83 (0.05)	1.0 (0.00)	1.0 (0.00)
Enron	0.87 (0.12)	0.76 (0.09)	0.96 (0.00)	0.96 (0.01)
Twitter	0.61 (0.18)	0.48 (0.14)	0.96 (0.07)	0.93 (0.05)

resents an unenforced, socially-adopted, expected behaviour in the agent population. Adherence to this is beneficial to all agents as it minimises clashing actions.

Agent interactions are often limited to neighbours in an underlying topology and convention emergence in static and dynamic topologies has been explored [3, 10, 11, 17]. Fixed strategy (FS) agents, who always choose the same action regardless of others' behaviour, have been shown to speed up the emergence when placed at topologically influential locations [4, 5, 15]. However, previous work assumes that the topology is fully observable to find these locations [6, 14, 15, 16]. This may not be possible for very large networks or those where access is rate-limited such as for Facebook or Twitter. Exploration of partially observable graphs has been performed as part of graph theoretic analysis [1, 2] and the influence maximisation problem [13] but most do not work under the assumption of finite, limited observations.

In this paper we investigate the effects of partial observability on FS agent placement in convention emergence. We introduce an algorithm, PO-PLACE, that finds influential locations given limited observations of the topology. We show that, with only 5-10% observation, PO-PLACE is able to give performance comparable to optimal degree placement.

2. EXPERIMENTAL SETUP

The partial observability problem can be described as any scenario where the network topology is initially unknown and is revealed incrementally by the use of *observations* of the neighbours of a given node. We define an observation as a retrieval of the neighbour list of a specific node, functionality that is often available in graph and online social network APIs. Given a finite number of these observations, o , we propose an algorithm, PO-PLACE, that attempts to find maximally high-degree nodes. When all observations have been used, PO-PLACE returns a list of the top n highest-degree nodes found.

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017)*, S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8-12, 2017, São Paulo, Brazil.
Copyright © 2017, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.




Figure 1: Comparison of PO-Place and Degree FS agent placement for convention emergence in real-world topologies. The y-axis indicates the proportion of runs where the desired strategy emerged as the convention.

PO-PLACE consists of two primary components: observation division/management and local area traversal. The former splits available observations as equally as possible amongst the number of starting locations specified, s , and reallocates unused observations from local area traversals that reach local optima. By having multiple such locations we can minimise the risk of all observations being allocated to a local area that results in suboptimal influential nodes. Additionally, as the locations are to be used for FS agents, multiple starting locations aids in separating out the FS agents, minimising redundant influence between them.

The second component deals with local area traversal. It maintains a priority queue of unexplored neighbours by degree and, as long as observations are still available, expands the top candidate to find the degrees of all its neighbours before iterating through this process. Thus it maintains a growing fringe of candidates which helps it avoid local maxima that approaches such as hill climbing may reach.

We make use of three real-world networks from the Stanford SNAP datasets [8]: CA-CondMat [7], the collaboration network of the arXiv COND-MAT (Condensed Matter Physics) category; Email-Enron [9], the email communications between workers at Enron; and Ego-Twitter [12], a crawl of Twitter follow relationships from public sources (for our purposes we ignore the directed nature of the edges). We varied both the number of nodes ($n = 5$ to $n = 30$) being requested as well as the number of observations provided ($o = 500$ to $o = 5000$ [$o = 3500$ for CondMat]). To establish an upper bound and allow comparison a full-observability degree placement was also performed for each of the networks with the same range of values. Each set of parameters was averaged over 30 runs.

We then used PO-PLACE to choose locations of FS agents and examined the convention emergence. For this, a population of agents is situated in the topologies introduced above and interactions are modelled as a game as in [10]. These were each run 50 times and the proportion of runs that emerged to the desired convention were found.

3. RESULTS AND DISCUSSION

Table 1 provides a summary of the extremes of the data with the standard deviations included in brackets. The right column is $o = 3500$ for CondMat and $o = 5000$ for Enron and Twitter. As can be seen, all networks respond well, even

with minimal numbers of observations. Even at $o = 500$, the degree sum of the nodes selected by PO-PLACE is often a substantial proportion of the optimal one. The performance varies across the three networks, with placement in CondMat doing best whilst Twitter is noticeably worse. This is to be expected, as 500 observations represents a substantially smaller proportion of the population in Twitter than it does in CondMat or Enron (0.61%, 2.34% and 1.48% respectively). Performance rapidly increases with the number of observations. For $n = 30$, the worst performing value of n , in both CondMat and Twitter $P_{deg-sum}$ exceeds 90% at round 5% network observation ($o = 1000$ for CondMat and $o = 5000$ for Twitter) and Enron exceeds 90% at around 10% observation ($o = 3500$). This shows that PO-PLACE is able to replicate the collective degree of pure degree placement with small proportions of observations. Varying s was shown to increase the effectiveness in many of the topologies.

We then used PO-PLACE with the best parameter combinations found previously to place FS agents within the system to influence convention emergence. These results are presented in Figure 1. As can be seen, in all networks, PO-PLACE is able to perform comparably to pure degree placement, despite only observing small fractions of the network. In particular, we show that Enron and Twitter respond well to increases in the numbers of observations available to them and that, with less than 10% of the network observed, they offer results only marginally worse than those requiring full observation. Increasing the number of starting locations has a substantial effect in the results from the Enron graph, dramatically increasing the effectiveness of PO-PLACE. It also has beneficial effects in the Twitter network, although this is less substantial and only at higher values of n .

4. CONCLUSION

We have presented an algorithm, PO-PLACE, to investigate and address the placement of fixed strategy agents for convention emergence in partially observable topologies. We showed that with observation of $<10\%$ of the network PO-PLACE can find locations with similar reach and influence as full-observability degree placement. We found that increasing the number of concurrent searches aids this capability in certain network conditions and that using these locations to place fixed strategy agents can influence convention emergence to similar levels as placing by degree.

REFERENCES

- [1] C. Borgs, M. Brautbar, J. Chayes, S. Khanna, and B. Lucier. The power of local information in social networks. In P. W. Goldberg, editor, *Proc. of the 8th International Workshop on Internet and Network Economics*, pages 406–419, 2012.
- [2] M. Brautbar and M. J. Kearns. Local algorithms for finding interesting individuals in large networks. In *Proc. of Innovation in Computer Science*, pages 188–199, 2010.
- [3] J. Delgado, J. M. Pujol, and R. Sangüesa. Emergence of coordination in scale-free networks. *Web Intelligence and Agent Systems*, 1(2):131–138, Apr. 2003.
- [4] H. Franks, N. Griffiths, and A. Jhumka. Manipulating convention emergence using influencer agents. *Autonomous Agents and Multi-Agent Systems*, 26(3):315–353, May 2013.
- [5] N. Griffiths and S. S. Anand. The impact of social placement of non-learning agents on convention emergence. In *Proc. of the 11th International Conference on Autonomous Agents and Multiagent Systems*, pages 1367–1368, 2012.
- [6] J. Kittock. Emergent conventions and the structure of multi-agent systems. In *Lectures in Complex Systems: Proc. of the 1993 Complex Systems Summer School*, pages 507–521, 1995.
- [7] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1(1):2, 2007.
- [8] J. Leskovec and A. Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.
- [9] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney. Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Mathematics*, 6(1):29–123, 2009.
- [10] J. Marchant and N. Griffiths. Manipulating conventions in a particle-based topology. In V. Dignum, P. Noriega, M. Sensoy, and J. S. Sichman, editors, *Coordination, Organizations, Institutions, and Norms in Agent Systems XI*, volume 9628 of *LNCS*, pages 242–261, 2016.
- [11] J. Marchant, N. Griffiths, and M. Leeke. Convention emergence and influence in dynamic topologies. In *Proc. of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1785–1786, 2015.
- [12] J. J. McAuley and J. Leskovec. Learning to discover social circles in ego networks. In *Neural Information Processing Systems*, pages 548–56, 2012.
- [13] S. Mihara, S. Tsugawa, and H. Ohsaki. Influence maximization problem for unknown social networks. In *Proc. of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 1539–1546, 2015.
- [14] N. Salazar, J. A. Rodriguez-Aguilar, and J. L. Arcos. Robust coordination in large convention spaces. *AI Communications*, 23(4):357–372, 2010.
- [15] S. Sen and S. Airiau. Emergence of norms through social learning. In *Proc. of the 20th International Joint Conference on Artificial Intelligence*, pages 1507–1512, 2007.
- [16] D. Villatoro, S. Sen, and J. Sabater-Mir. Topology and memory effect on convention emergence. In *Proc. of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology*, pages 233–240, 2009.
- [17] A. Walker and M. Wooldridge. Understanding the emergence of conventions in multi-agent systems. In *International Conference on Multi-Agent Systems*, pages 384–389, 1995.