

Convention Emergence in Partially Observable Topologies

James Marchant and Nathan Griffiths

Department of Computer Science, University of Warwick
{james, 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, whose interactions are constrained by an underlying network, 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

Coordinating the actions of independent agents within a multi-agent system (MAS) increases efficiency within the system. Incompatible action choices made during interactions can cause clashes, which may incur resource costs and limit the overall effectiveness of the system. Establishing protocols of interaction, such as which action to choose in a given situation, minimises such clashes and helps to maximise the potential of the system.

However, it is not always possible to dictate such rules and protocols in a top-down manner. In multi-agent systems, with agents controlled by multiple parties or systems which lack a centralised control mechanism, it is often infeasible to establish this level of *a priori* coordination. Additionally, for systems where the range of choices available to agents is large or has no evident optimal selection, it may be undesirable to enforce rules of this nature.

Convention emergence allows a system to deal with these problems in a decentralised, online manner. A *convention* represents a socially-adopted expected behaviour amongst agents, for instance the correct course of action in a given scenario. Convention emergence has been shown to be possible in both static and dynamic networks with minimal requirements, needing only rational agents that are able to learn [6, 13, 22].

Fixed strategy (FS) agents are those that continue to choose the same action regardless of the behaviour of others around them or the results of their actions. Placing such agents within a system has been shown to affect the direction and speed of convention emergence, with small numbers of FS agents eliciting change in much larger populations [19]. In systems constrained by an underlying network topology, placing such agents by heuristics based on network features such as degree magnifies this effect [8, 9].

Previous work on convention emergence often assumes that the topology constraining agent interactions is fully observable, allowing highly influential locations to be found easily [11, 17, 19, 21]. However, in many real-world applications such information is not always readily available. This can be due to factors such as the problem size or external limitations such as restricted access to network information or a network’s API as is the case with Twitter or Facebook.

In this paper we explore the effect of the restrictions placed on FS agent placement in partially observable topologies. We propose an algorithm, PO-PLACE, to find influential locations within such topologies given a highly limited number of network queries. We show the effectiveness of PO-PLACE at finding approximations of the highest degree locations for several real-world topologies under a number of restrictions on available information. We then apply PO-PLACE to select FS agents within these networks and examine the effect on convention emergence compared to placing with full topological knowledge. This approach allows an interested third party, with limited access to the system, to find the appropriate locations to target their influence efforts.

The remainder of this paper is arranged as follows. In Section 2 we explore related work on convention emergence and local information strategies for finding influential nodes. Section 3 describes the algorithm and its design, whilst Section 4 describes the network datasets and experimental setup. Section 5 contains the analysis and discussion of the results and Section 6 concludes the paper.

2 Related Work

Ensuring coordination in MAS allows increased system efficiency and conventions are a lightweight method of doing so. Conventions place ‘soft constraints’ on agent choices by encouraging mutually beneficial behaviour by adherence to the convention. Unlike *norms* there is no explicit punishment for going against the convention but doing so is likely to incur a cost to the agent due to increased clashes often represented as a negative interaction payoff [10, 18]. Conventions can thus be described as “an equilibrium everyone expects in interactions that have more than one equilibrium” [24]. Agents adhering to the convention expect

others to behave in a certain way and, because of this, can act efficiently when this expectation is met. Conventions have been shown to emerge unaided from local agent interactions in systems [6, 11, 20, 22] and require no additional or assumed agent capabilities to enable punishment (as is the case for norms). The only assumptions necessary for conventions to emerge are that agents are rational and have the capability to learn from their interactions. Numerous works have shown that rapid and robust convention emergence occurs with these minimal assumptions [9, 19, 22].

‘Social learning’ has been proposed as a way for agents to converge on a convention where agents monitor payoffs they receive from their choices when interacting with others and use a simplified Q-Learning algorithm to inform future decisions [19]. The payoffs directly quantify the notion of an action clash costing resources and convention emergence can occur without explicit memory of the interaction. However, the work does not consider a connecting topology that limits agent interactions. In many application domains such a topology is likely, whether it be a social network or a more explicit communication network and can have a large effect on the nature of convention emergence [5, 21].

Despite lacking a connecting topology, Sen and Airiau’s work introduces the concept of fixed strategy (FS) agents, those agents which always choose the same action regardless of the current situation or convention, as a way to influence convention emergence. They show that a small number of such agents is able to manipulate the convention emergence within a much larger population. Griffiths and Anand [9] expand on this by considering FS agents in a network topology. In their model, all agents are situated as nodes within the network and interactions are limited to neighbours. They showed that *where* FS agents are placed is a key factor in their effectiveness. Placing the FS agents at influential locations such as nodes of high degree or betweenness centrality offers substantially better performance than random placement. This was explored further by Franks et al. [7, 8] who included more advanced placement metrics such as eigencentrality.

This previous work assumes full visibility of the network topology to inform FS placement. Indeed, little work on partial observability for convention emergence has been done. This paper expands the state of the art by considering the effect of restricted observations on the ability to robustly and efficiently place FS agents in static, real-world topologies. Convention destabilisation [14] and dynamic topologies [13] will be investigated in future work.

Related work exists in the fields of graph algorithms and influence spread, the latter sharing many qualities with convention emergence. For instance, Brautbar and Kearns present a novel model [2], *Jump and Crawl*, motivated by operations commonly available in networks such as Facebook. Their model consists of two aspects: *Jump* which moves to a randomly selected node in the network and *Crawl* which searches all neighbours of the selected node for high-degree nodes. They provide bounds for many different types of network but, for an arbitrary network, finding the highest degree node approaches $O(n \log n)$, a large factor for even medium-sized networks.

The influence maximisation problem [3, 4] attempts to find a selection of nodes such that the spread of influence (often modelled as single chance ‘cascades’) from them is maximised. As in this paper, Mihara *et al.* [15] assume the network is initially unknown and show that influence maximisation effectiveness of 60-90% with 1-10% network observation is achievable. This work also uses a ‘growing fringe’ approach with priority based on degree estimation. As influence maximisation and convention emergence are similar in aim, this indicates that results are achievable under partial observability constraints.

Whilst many of these approaches are similar in application they differ in that our investigation focuses on the often encountered scenario of limited, finite observations. Making optimal use of these is paramount and so necessitates a different set of considerations.

3 Placement Strategy

In this paper, the partial observability problem for networks can be described as any scenario where a network’s topology is initially unknown and is revealed incrementally within a local neighbourhood of nodes already explored [1]. As a solution to the partial observability problem for FS agent selection we propose a heuristic algorithm, PO-PLACE. This section describes the function of the algorithm as well as the justification for the design choices.

3.1 Partial Observability Algorithm

The placement strategy is presented in Algorithms 1 and 2 and has the following aim: Given a network, $G = (V, E)$, a desired number of locations, n , and a limited number of observations, o , find a selection of nodes $\{v_1, \dots, v_n\} \subset V$ such that *degsum* = $deg(v_1) + \dots + deg(v_n)$ is maximised. We define an observation as a query that retrieves the list of neighbours, $N(u)$ for a given node, u . This functionality is frequently available in real-world network APIs (such as Twitter or Facebook) and so we assume that such information is available. This assumption is later relaxed to allow the algorithm to explore situations with only limited neighbour information. We assume that the set of nodes, V , is known but the set of edges, E , (and hence neighbours and degree of a node) is not. Finding the highest degree nodes is desirable since fixed strategy agent placement by degree consistently produces effective convention emergence [8, 9, 13, 14] but without requiring computationally expensive metrics such as betweenness centrality. The degree of nodes can be entirely derived from local information and, as such, is an applicable heuristic within partially observable networks.

The algorithm begins by creating an empty set, S , to monitor which nodes have already been explored and an empty mapping, N , that maps a node v to $N(v)$, its set of neighbours. As we only consider static topologies in this paper, by storing this information we can avoid using observations redundantly.

Many of the other approaches [1, 15] to finding high-degree nodes select a random starting node and then ‘grow’ outwards, selecting the highest degree

Algorithm 1 Partial Observability Placement

```
1: procedure PO-PLACE( $G, n, o, s, p, f$ )
2:   Create empty node set,  $S$ 
3:   Create empty mapping,  $N$ 
4:    $o_{rem} \leftarrow o$ 

5:   while  $o_{rem} > 0 \wedge |S| < |V|$  do
6:     Select  $v$  u.a.r from  $\{V \setminus S\}$ 
7:     if  $o_{rem} \bmod s \neq 0$  then
8:        $o_{local} \leftarrow \min(\lceil o/s \rceil, o_{rem})$ 
9:     else
10:       $o_{local} \leftarrow \min(\lfloor o/s \rfloor, o_{rem})$ 
11:    end if
12:     $o_{rem} \leftarrow o_{rem} - o_{local}$ 
13:     $o_{unused} \leftarrow \text{Traverse}(G, o_{local}, v, p, f, S, N)$ 
14:     $o_{rem} \leftarrow o_{rem} + o_{unused}$ 
15:  end while

16:  return  $n$  highest-degree nodes in  $S$ 
17: end procedure
```

nodes from the neighbourhood surrounding those already explored. However, this is not desirable in FS agent placement since, with limited observations, it is likely to produce a single cluster of well-explored nodes. Selecting from this cluster will then mean that all FS agents are close together, making some of their influence redundant. Instead, we build on the notion of *Jump and Crawl* [2]. We explore a local area up to a defined amount and then ‘jump’ to another location and explore around this new point. This helps to minimise the risk of overlap between high-degree nodes, as well as ensuring that a bad initial random selection does not hinder the final selection.

To facilitate this, we introduce a parameter, s , which dictates the minimum number of separate local area explorations that will take place. The observations are split, as evenly as possible, between each of these explorations with the earlier ones receiving any spare observations (this is achieved between Lines 7 and 11 of Algorithm 1). This subset of observations is then passed to the local area traversal which is presented in Algorithm 2. If any observations are unused by the local area traversal (for instance if it finds a local maxima) they are returned to the pool of available observations and used in later, additional local traversals.

Algorithm 2, TRAVERSE, describes the local area traversals. It is passed both S and N , to avoid redundant exploration, as well as the initial start node of the local area, v . It is also passed its own local limit of observations and two parameters from outside, p and f , which are explained below. It maintains a max-priority queue to determine which node(s) it should next explore by highest degree and begins by adding v to this queue. Throughout Algorithm 2, observation of a node’s neighbour list is stored in N to avoid additional queries. The algorithm then performs the following, until either the queue is empty or all assigned observations have been used up:

1. Take the top f nodes from the queue (or all elements, if fewer).
2. For each of these nodes, find the set of unexplored nodes in its neighbours.

Algorithm 2 Local area traversal algorithm

```
1: procedure TRAVERSE( $G, o, v, p, f, S, N$ )
2:   Create max-priority queue,  $Q$ 
3:    $count \leftarrow 0$ 
4:   if  $v$  not in  $N$  then
5:      $N[v] \leftarrow N(v)$ 
6:     Add  $v$  to  $S$ 
7:      $count \leftarrow count + 1$ 
8:   end if
9:   Add  $(v, |N[v]|)$  to  $Q$ 

10:  while  $|Q| > 0 \wedge count < o$  do
11:     $Fringe \leftarrow$  top  $min(f, |Q|)$  elements of  $Q$ 
12:    for all  $u$  in  $Fringe$  do
13:       $Avail \leftarrow \{N[u] \setminus S\}$ 
14:       $num \leftarrow min(|Avail|, max(f, \lfloor p \times |Avail| \rfloor))$ 
15:       $Chosen \leftarrow$  u.a.r select  $num$  members of  $Avail$ 
16:      for all  $w$  in  $Chosen$  do
17:         $N[w] \leftarrow N(w)$ 
18:        Add  $w$  to  $S$ 
19:         $count \leftarrow count + 1$ 
20:        if  $count = o$  then
21:          return 0
22:        end if
23:        Add  $(w, |N[w]|)$  to  $Q$ 
24:      end for
25:    end for
26:  end while

27:  return  $o - count$ 
28: end procedure
```

3. Choose a proportion, p , of these (or up to f if this proportion would be less than f).
4. Add these nodes to the queue after finding their neighbours.

Parameter f is the ‘fringe size’, the number of nodes that are expanded simultaneously before their neighbours are queued. This acts as a control over how ‘breadth-first’ or ‘depth-first’ the local traversal approach will be. Parameter p is the proportion of the node’s neighbours that should be queried. This allows the algorithm to simulate situations where a node’s full neighbour list is either not fully available (for instance, an API that only returns a subset) or where doing so incurs additional cost. In the latter case we seek to explore the effect that only querying p proportion of neighbours has on the performance of PO-PLACE. Whilst it will reduce the effectiveness, establishing the extent of this reduction, and whether the results are still close enough to degree placement, allows PO-PLACE to be effective over a wider range of scenarios.

4 Experimental Setup

This section defines the real-world topologies and the experimental setup used for analysis of PO-PLACE. We then describe the model of convention emergence used to study the efficacy of PO-PLACE for FS placement selection.

4.1 Networks

We make use of three real-world networks from the Stanford SNAP datasets [12]. These datasets represent a number of different methods of social interaction and, as such, each have different features allowing a wide-ranging look at the effectiveness of PO-PLACE. The three datasets chosen are: CA-CondMat, the collaboration network of the arXiv COND-MAT (Condensed Matter Physics) category; Email-Enron, the email communications between workers at Enron; and Ego-Twitter, a crawl of Twitter follow relationships from public sources (for our purposes we ignore the directed nature of the edges). These datasets are used frequently in both convention emergence and influence spread research [3, 7, 16, 23] as performance benchmarks.

For the purposes of monitoring convention emergence in these networks, we only want to examine a single, connected component. As such, all 3 networks were reduced to their largest weakly connected component (WCC). Additionally,

Table 1. Original and Modified Network Sizes

	Network		Largest WCC	
	$ V $	$ E $	$ V $	$ E $
CA-CondMat	23,133	93,497	21,363	91,286
Enron-Email	36,692	183,831	33,696	180,811
Twitter	81,306	1,768,149	81,306	1,342,296

any self-loops (edges from a node to itself) were removed as such edges artificially inflate a node’s degree whilst not increasing its ability to influence others. Table 1 shows the number of nodes and edges in each network and the number of nodes and edges (without self-loops) in their largest WCC.

4.2 Experimental Setup

We performed simulations of PO-PLACE on the real-world networks described above. 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.

For convention emergence, a population of agents is situated in the topologies. Each timestep, each agent chooses one of its neighbours u.a.r to play the 10-action coordination game [19] receiving positive or negative payoffs depending on whether their choices match. Agents use a simplified Q-Learning algorithm to learn the most beneficial choice. We utilise the 10-action game as used by [14] to avoid the issues of small convention spaces raised in Section 2 and to allow comparison to previous work. They have a chance to randomly choose their action ($p_{explore} = 0.25$) or else choose the most beneficial one. FS agents replace the agents at the chosen locations and always choose their predetermined action.

5 Results & Discussion

In this section we present the analysis of PO-PLACE and compare it to the upper bound from degree placement. We explore the effects of the various parameters




Fig. 1. $P_{deg-sum}$ and $deg-sum$ performance of PO-PLACE for varying n (# of locations) and o (# of observations) in the real-world networks.

on PO-PLACE at different levels of observation. We then use these findings as insight to compare the performance of PO-PLACE to degree for convention emergence when used to place FS agents into the chosen networks.

5.1 PO-PLACE Output

We begin by looking at the isolated algorithm output, comparing it to the output generated by a degree placement scheme. As the aim of PO-PLACE is to maximise $deg-sum$ this is our primary metric by which to evaluate PO-PLACE. The highest $deg-sum$ possible in each network is that of the set of highest degree nodes. Establishing this as an upper bound allows evaluating the performance of PO-PLACE by comparing the $deg-sum$ of its output as a proportion of that of the pure degree network. We denote this as $P_{deg-sum}$.

Whilst $deg-sum$ describes the maximum reach of the nodes selected, another useful metric is the size of the 1-hop neighbourhood of those nodes. This can be defined as: $1-HOP(L, G) = \{v \in \{V \setminus L\} | \exists(u, v) \in E \wedge u \in L\}$ where L is the set of nodes selected for placement and $G = (V, E)$ is the network. That is, the 1-Hop neighbourhood is the set of nodes that are connected to a member of S but are not in S themselves. The 1-Hop neighbourhood offers a slightly different measure of influence by discounting nodes that are connected to multiple members of S . Whilst normally tied closely to $deg-sum$ a noticeable disparity indicates that the selected nodes are likely to be clustered close to one another, which is undesir-

able. As with *deg-sum* we concern ourselves with the proportionate behaviour of 1-HOP size, $P_{|1\text{-HOP}|}$.

The final metric we use to evaluate the performance is based on the Jaccard Index which measures similarity between two sets. The Jaccard Index is defined as $J(A, B) = |A \cap B| / |A \cup B|$. However, in our instance, one of the sets is static. We are trying to approximate that set with the other (i.e. a one-way similarity), whilst the Jaccard Index is looking at the two-way similarity between them. Instead we want to measure how close the selection of PO-PLACE is to the baseline, and so we define a distance measure, D_{Base} , thus: $D_{Base}(L, Base) = |L \cap Base| / |Base|$. That is, the fraction that elements of L make up of the baseline set, $Base$. This metric enables evaluation of how close the actual node selection of PO-PLACE is to that of degree placement, whilst the previous two measure the selection’s features.

These metrics offer insight into the influence and reach of the nodes selected by PO-PLACE as well as allowing a direct comparison to degree-based placement with full observability. Thus they should be good predictors of the performance of PO-PLACE in the convention emergence setting.

Varying Observations We begin by considering the base case of the algorithm where $s = p = f = 1$. This allows us to study the effect of varying the number of observations and provides a lower bound on the expected performance of PO-PLACE. With these settings, PO-PLACE closely resembles the algorithms presented by Borgs *et al.* [1] and Mihara *et al.* [15].

We examine the effects of varying both the number of observations available (o) as well as the number of locations requested (n) in all three networks. For all networks, n was varied between 5 and 30 in increments of 5 and o was varied from 500 observations up to 3500 (for CondMat) or 5000 (for Enron and Twitter). The results are presented in Figure 1.

As can be seen in Figure 1, 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 where it varies from 90% ($\pm 5\%$) at $n = 5$ to 83% ($\pm 5\%$) at $n = 30$. The algorithm similarly performs well in Enron, though to a lesser extent. The performance in Twitter is noticeably worse, varying from 61% to 48% with larger standard deviations for both. 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). Even with this, the percentage achieved in Twitter with such limitations substantially outperforms the naïve solution of using all observations at random locations (16% ($\pm 6\%$) for $n = 5$, $o = 500$, averaged over 100 runs).

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 ($o = 3500$).

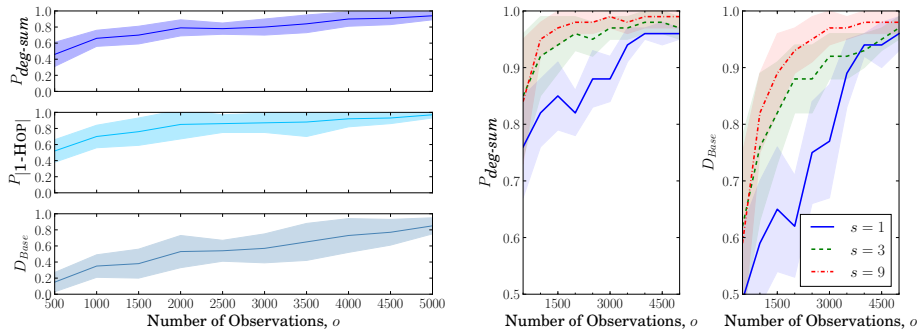


Fig. 2. Metric performances of PO-PLACE for the Twitter network, $n = 20$. **Fig. 3.** Effect of varying s on $P_{deg-sum}$ and D_{Base} . Enron network, $n = 30$.

Figure 1 also shows that the relationship between $P_{deg-sum}$ and increasing o is one of diminishing returns, with improvements in $P_{deg-sum}$ most noticeable at lower values of o . This is to be expected, the relative increase in o is smaller at higher values, but dictates that increasing the effectiveness of PO-PLACE at low values of o will have the most benefit. Additionally, in each network, the difference in performance across the values of n becomes less noticeable at higher o . Thus, any increased performance from PO-PLACE will be most noticeable early on.

The other metrics we use to evaluate PO-PLACE show similar behaviour to $P_{deg-sum}$, increasing rapidly with the number of observations. Figure 2 shows a representative example of the three metrics' variation with o for the Twitter network when requesting 20 locations. The shaded regions represent the standard deviations. As can be seen, both the $deg-sum$ and 1-HOP proportions increase rapidly up until $o = 2000$ and then any further gains occur over longer spans. The standard deviations for each of these decrease as well, from approximately 15% at $o = 500$ down to around 5% at $o = 5000$. This indicates that, not only is PO-PLACE finding sets of nodes with higher degree, it is doing so consistently at higher numbers of observations, a finding that is repeated across all networks and values of n . $P_{|1-HOP|}$ is consistently at the same level, if not better than, $P_{deg-sum}$. Whilst the two should be well-correlated, this shows that PO-PLACE is not simply choosing nodes close to one another and, indeed, is often choosing nodes that have a better neighbourhood size than the $deg-sum$ would indicate.

The performance of PO-PLACE when evaluated by D_{Base} is noticeably different than the other two metrics and offers an interesting insight. The same pattern of diminishing returns is not present and D_{Base} continues to increase with additional observations. Note that, although both the degree sum and neighbourhood size are comparable to that of pure degree placement, the low values of D_{Base} indicate that the nodes selected are not the same as the actual highest degree nodes. Section 5.2 evaluates whether this difference has a noticeable effect on convention emergence or if the reach and influence indicated by high $deg-sum$ and 1-HOP scores is the best indicator of success as hypothesised.

Varying Concurrent Searches Having established a baseline for PO-PLACE and explored the effects of limited observations we now explore the variants of the algorithm. As noted in the prior section, at low values of o the *deg-sum* performance of PO-PLACE is consistently lower, with performance in the Twitter network as low as 48%. With very limited observations, making the best use of them is paramount. In Section 3 we hypothesised that splitting the available observations between multiple locations in the network and exploring them in parallel may offer improvements over crawling from a singular location.

To test this hypothesis, we varied s from 1 to 9 to determine the effect that these concurrent searches would have. Figure 3 shows a typical case in the Enron network for $n = 30$. Shaded areas represent the errors of each plot. The left-hand graph shows the effect on $P_{deg-sum}$ of varying the number of concurrent searches, splitting the observations between them. As can be seen, adding concurrent starting points has an immediate and noticeable effect, especially at low numbers of observations. At $o = 500$ the proportion achieved by *deg-sum* is 10% higher when additional starting locations are introduced and this difference becomes even more noticeable as o increases. Indeed, for most values of o , adding additional starting locations had significant benefits in both the Enron and Twitter networks, with the benefits become less marked at high o where $P_{deg-sum}$ approaches 1.0 unaided. Whilst there is a noticeable drop-off in effectiveness after initial parallelisation ($s = 5$ and $s = 7$, not included in the results to aid readability, offer little improvement over $s = 3$ for example) the effect at low values of s is substantial as can be seen. Concurrent starting points enable saturation of the algorithm’s effectiveness at much lower values of o and not only increase $P_{deg-sum}$ and $P_{|1-HOP|}$ (not pictured) but, as shown in Figure 3, cause marked improvement in D_{Base} as well, indicating that this change facilitates much better approximation of the degree placement.

However, it should be noted that this pattern is not consistent. In the CondMat network, increasing s had little effect and in a few settings was actually detrimental. This indicates that there is perhaps an underlying feature of the CondMat topology that benefits from localised crawling and will be an area of future study. The results of CondMat in Figure 1a lend additional weight to this hypothesis, with behaviour that is substantially different than the other two topologies despite being of comparable size to Enron. Overall though, increasing s by even a small amount is likely to benefit the performance of PO-PLACE.

Partial Neighbour Lists In many settings, retrieving the whole of an agent’s neighbour list may also be impossible. Whether this is due to a technical limitation (only being able to retrieve a certain percentage of information) or because such information is not publicly available and is instead reserved for ‘premium’ or ‘subscribed’ users of such a network, ensuring that PO-PLACE is robust to such issues is a necessity to make it widely viable.

To simulate these restrictions, and measure their effect on the performance of PO-PLACE, the parameter, p , controls the proportion of an agent’s neighbours that may be explored. Results until this point have assumed that the full neigh-




Fig. 4. Effect of varying p on $P_{deg-sum}$. Twitter network, $n = 5$.




Fig. 5. Effect of varying f on $P_{deg-sum}$. Enron network, $n = 30$.

hour list for any agent is available upon request (i.e. $p = 1.0$). p is varied between 0.3 and 0.9 to determine the impact of this limitation. Representative results are shown in Figure 4 for the Twitter network and $n = 5$ but are applicable across all networks and values of o and n .

The results in Figure 4 show that different values of p have minimal effect on the performance of PO-PLACE. For all values of p , $P_{deg-sum}$ is comparable. Performing a 95% confidence interval Welch’s t-test against the $p = 1.0$ results at each point, only $p = 0.3$ ($o = 1500, 2000, 3500$) and $p = 0.5$ ($o = 1500, 3500$) are significantly worse. This pattern of minimal difference is repeated in all networks, with none seemingly more susceptible or affected by partial neighbour lists. We conclude that PO-PLACE is robust to receiving only partial information of this nature and is primarily unaffected by such limitations.

Breadth-First vs Depth-First Expansion Finally, we turn our attention to the concept of breadth-first vs depth-first expansion in PO-PLACE. That is, when crawling the local area, should additional current area expansion be performed before considering new additions (breadth-first) or purely iteratively (depth-first). Where there is locally a clearly defined degree gradient we expect the latter to perform better. However, depth-first expansion also risks expending all the observations whilst exploring a suboptimal, locally maximal path.

Parameter f allows study of this by controlling how many of the current highest degree nodes that PO-PLACE is aware of are expanded concurrently. Experiments up until now have had $f = 1$ (depth-first). We now vary f from 1 to 9. Figure 5 presents these findings in the Enron network for $n = 30$. As with the previous results, it is our finding that the patterns here are replicated throughout the different topologies and values of n .

Similar to the findings when varying p , varying f has little absolute impact on the capabilities of PO-PLACE. However, using a 95% confidence interval Welch’s t-test, all but $f = 9$ are statistically significantly worse at $o = 500$. This is likely due to the limited observations being focused too locally. All are significantly better between $o = 2000$ and $o = 3000$ but there is little gain in selecting values of f beyond 3 as the performance of PO-PLACE is almost identical. Overall,

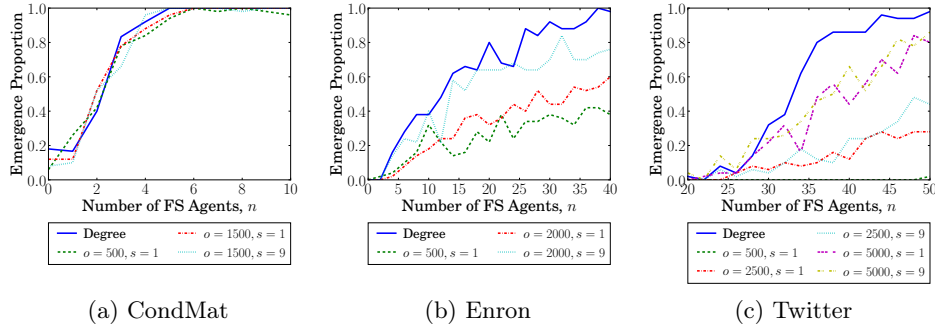


Fig. 6. Comparison of PO-PLACE and Degree FS agent placement for convention emergence in real-world topologies. The y-axis indicates the proportion of runs where the desired strategy emerged as the convention.

PO-PLACE seems to gain little from considering the local area more thoroughly before further expansion. Whether this is intrinsic in the design or a facet of the topologies being explored is ongoing work.

5.2 Convention Emergence under Partial Observability

Having explored the performance of PO-PLACE under different topologies and types of partial observability, we now examine how PO-PLACE compares to degree placement for FS agents in convention emergence in static networks. Having established ranges of parameters that offer the best performance improvements for each topology, these will be utilised to compare the algorithm to degree placement. Additionally, basic settings (small numbers of observations, no concurrent placements) provide a baseline comparison of PO-PLACE.

A convention has emerged when the population has converged to have one action as the dominant choice of agents in the network. Most work considers this to be the case when the convention reaches 90% dominance [9, 13, 19]. However, much of this work utilises synthetic networks rather than real-world topologies and populations that are substantially smaller than those we consider. Preliminary experiments show that the topologies are relatively resistant to convention emergence, requiring both high numbers of FS agents as well as substantial time. As we are concerned with a comparison of the performance of PO-PLACE against pure degree placement we wish to find settings that are guaranteed to repeatably experience convention emergence. As such, we consider a convention to have emerged when the 80% Kittock criteria is met, $K_{80\%}$ [11]. That is, a convention has emerged when 80% of the population, when not exploring, would choose the same action. This indicates a high level of dominance of the desired action and allows more robust comparisons. We find that such a threshold is reliably reached, if it is likely to be reached at all, within 10000 iterations for the CondMat and Twitter networks and within 15000 iterations for the Enron

network. As such, we measure the proportion of runs that have converged to the desired strategy within these time-frames across all networks.

The results are presented in Figure 6. We utilise PO-PLACE and Degree Placement to allocate FS agents across a range found to exhibit noticeable changes in convention emergence rates with the parameters indicated. The values of o chosen within each topology are such that the number of observations is, at most, approximately 5% of the agent population. All runs are performed 50 times and the proportion of runs that produce the desired convention (strategy chosen u.a.r at time $t = 0$ and assigned to all FS agents) is measured.

Figure 6a shows the results for the CondMat topology. As was expected, due to the behaviour of CondMat in the PO-PLACE experimentation, all of the chosen parameters produce comparable results to the pure degree placement. Even at the worst performing parameters ($o = 500, s = 1$) there is no discernible difference between the performance of degree placement and PO-PLACE, whilst at higher number of observations (where PO-PLACE was entirely approximating the highest-degree nodes as seen in Figure 1a) the performance is as expected. Of note is the fact that, whilst it resulted in worse output of PO-PLACE in the prior section, increasing s does not noticeably affect the performance here.

Within the other networks the difference in performance is more noticeable but still indicates that PO-PLACE is generating close approximation of the degree placement. In both Enron and Twitter (Figures 6b and 6c) the minimal observation situation performs substantially worse than degree placement, particularly in the Twitter network. However, when given observations of 5% of the network, PO-PLACE performs noticeably better. Whilst it still falls behind the performance of degree placement in both networks the difference is substantially smaller with PO-PLACE performing around 50-70% as effectively on average as degree placement in both networks (0.52 ± 0.08 in Enron, 0.69 ± 0.18 in Twitter). However, when we increase s , as was found in Section 5.1, it improves this substantially to 0.82 ± 0.15 average effectiveness compared to degree placement in Enron and, less substantially, to 0.79 ± 0.3 in the Twitter network. We quantify these values by comparing the emergence proportions of PO-PLACE and degree at each value of n and calculating the ratio between them which we then average. We discount values where either placement is achieving less than a 0.1 emergence proportion to avoid noisy results influencing the measure. As 0.1 is the expected emergence proportion of our desired strategy in a convention emergence we do not influence, we believe discounting values below this allows a more accurate comparison between the two algorithms. In the Twitter network, we also consider $o = 2500$ as the effect of increased s was more pronounced for this value during Section 5.1. Whilst there is a noticeable improvement at higher n the average compared effectiveness differs only marginally: 0.24 ± 0.06 for $s = 1$ and 0.3 ± 0.11 for $s = 9$. In the Twitter network, o is the dominant factor.

Overall, we have shown that even when only observing a small portion of the underlying topology, and strategically using these observations to maximise their effect, it is possible to achieve comparable performance to degree placement with full network visibility using PO-PLACE.

6 Conclusion

Finding influential positions within a network topology to maximise the effectiveness of fixed strategy agents is an ongoing area of research in convention emergence. The problem has many facets and variations that make it difficult to find an optimal yet general approach. In many cases, placing the fixed strategy agents at high degree nodes provides effective convention emergence with little computational overhead. Finding high-degree nodes in a network is trivial when the network is fully observable. In many domains, this may not always be possible. Technical limitations such as memory constraints or incomplete information and usage limitations such as finite API calls mean that often a network topology may only be *partially observable*. Finding effective placement for FS agents with these restrictions adds another level of complexity.

In this paper we presented a placement algorithm, PO-PLACE, that is designed for use in partially observable topologies. It uses finite observations to find sets of high-degree nodes and approximates the set of nodes that would be selected given full observability.

With small proportions of the network being observable, PO-PLACE can locate nodes with similar reach and influence as degree placement. We evaluate the performance in three real-world topologies and show that the addition of concurrent searches and splitting of observations improves the performance of the algorithm across all metrics. With 1-10% observation the algorithm is able to find sets of nodes with >90% of the reach and influence of degree placement.

Finally, we showed that PO-PLACE performs comparably to degree placement when used to facilitate convention emergence using fixed strategy agents whilst only observing 5% of a network topology. We found that the additional aspects of PO-PLACE benefit the placement mechanism and demonstrated that convention emergence is easily facilitated in partially observable networks.

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] D.-B. Chen, R. Xiao, and A. Zeng. Predicting the evolution of spreading on complex networks. *Scientific Reports*, 4:6108, 2014.
- [4] W. Chen, Y. Wang, and S. Yang. Efficient influence maximization in social networks. In *Proc. of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 199–208, 2009.
- [5] J. Delgado. Emergence of social conventions in complex networks. *Artificial Intelligence*, 141(1–2):171–185, 2002.
- [6] 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.

- [7] H. Franks, N. Griffiths, and S. S. Anand. Learning agent influence in MAS with complex social networks. *Autonomous Agents and Multi-Agent Systems*, 28(5):836–866, Sept. 2014.
- [8] 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.
- [9] 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.
- [10] M. Kandori. Social norms and community enforcement. *The Review of Economic Studies*, 59(1):63–80, 1992.
- [11] 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.
- [12] J. Leskovec and A. Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.
- [13] 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.
- [14] J. Marchant, N. Griffiths, M. Leeke, and H. Franks. Destabilising conventions using temporary interventions. In *Coordination, Organizations, Institutions, and Norms in Agent Systems X*, volume 9372 of *LNCS*, pages 148–163, 2015.
- [15] 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.
- [16] S. Pei, S. Tang, and Z. Zheng. Detecting the influence of spreading in social networks with excitable sensor networks. *PloS One*, 10(5):e0124848, 2015.
- [17] N. Salazar, J. A. Rodriguez-Aguilar, and J. L. Arcos. Robust coordination in large convention spaces. *AI Communications*, 23(4):357–372, 2010.
- [18] B. T. R. Savarimuthu, R. Arulanandam, and M. Purvis. Aspects of active norm learning and the effect of lying on norm emergence in agent societies. In D. Kinny, J. Y.-j. Hsu, G. Governatori, and A. K. Ghose, editors, *Agents in Principle, Agents in Practice*, volume 7047 of *LNCS*, pages 36–50. Springer, 2011.
- [19] 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.
- [20] D. Villatoro, J. Sabater-Mir, and S. Sen. Social instruments for robust convention emergence. In *Proc. of the 22nd International Joint Conference on Artificial Intelligence*, pages 420–425, 2011.
- [21] 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.
- [22] 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.
- [23] X. Wang, Y. Su, C. Zhao, and D. Yi. Effective identification of multiple influential spreaders by DegreePunishment. *Physica A: Statistical Mechanics and its Applications*, 461:238–247, 2016.
- [24] H. P. Young. The economics of convention. *The Journal of Economic Perspectives*, 10(2):105–122, 1996.