Spatial Community Structure Promotes Extremism in an Adaptive Spatial Social Network Model

Seth Bullock

School of Computer Science, University of Bristol, UK
seth.bullock@bristol.ac.uk

Abstract

A model of opinion dynamics amongst agents embedded on an adaptive social network is extended to introduce tuneable *spatial embedding*. As in the original model, the opinions and social connections of a population of model agents change due to three social processes: *conformity*, *homophily* and *neophily*. Here, however, direct interactions are constrained to take place only between pairs of agents that are linked by short spatial connections, or between pairs of agents that have benefited from some degree of random rewiring of these spatial connections. This introduction of spatiality could be expected to either *reduce* the ability of extreme agents to connect with one another in order to form extreme communities, or to *increase* their ability to influence the opinions of the community of agents clustered around them. Results demonstrate that the latter is the case. Spatial constraints tend to encourage extreme communities (relative to comparable non-spatial networks) due to the increased number and strength of distinct agent communities in spatial networks. These results suggest that the presence of strong community structure (rather than high clustering coefficients or short characteristic path lengths) may promote extremist communities in real-world populations.

Introduction

For most of evolutionary history, direct interactions between organisms have been strongly constrained by their spatial proximity to one another. During this period, the most far-reaching interactions might have included visible displays (Bullock and Cliff, 1997), audible calls (Bullock, 1998), and stigmergic sensory markers (Bullock, 2016) each of which can tend to lead to correlated systems due to the enabling constraints of low-dimensional spatial embedding (Bullock and Buckley, 2009; Bullock and Geard, 2010). However, with the advent of information technologies, modern humans have gained the ability to interact over much longer distances. This "death of distance" (Cairncross, 2001) has hugely increased our ability to spread information, opinions and ideas rapidly across the globe, thereby radically altering our capacities for large-scale co-ordination, co-operation and creative innovation on the one hand, but also competition, consumption, and thoughtless copying on the other.

More specifically, the tendency for online social interactions to spread ideas and influence opinions, and for online communities to organise around these ideas and opinions, is gaining attention in the context of what appears to be an ongoing period of increasing social polarisation (Levin et al., 2021). The costs of forming and maintaining connections on online social networks and of broadcasting information on these networks are low and are also largely uncorrelated to the physical distances involved. Has this contributed to the increasingly fragmented and polarised nature of online communities? If we are to address the echo-chamber effects of social network self-organisation (Axelrod et al., 2021; Santos et al., 2021) we will need to understand the interplay of spatial and non-spatial processes that underpin online opinion dynamics and community formation. Analogously, engineers interested in designing or managing distributed multi-agent systems may have similar interests in understanding flows of influence within populations of collaborating agents operating in either virtual or physical spaces (Pitonakova et al., 2016a,b, 2018; Jacyno et al., 2009, 2013).

Across these different settings, community formation can be understood as an ongoing reflexive process of coevolutionary adaptation, i.e., a social network constrains which agents are more likely to interact with one another and, simultaneously, is shaped and reshaped by these interactions (zu Erbach-Schoenberg et al., 2013; Sayama, 2020a). Models of adaptive social networks, where network topology and node traits co-evolve simultaneously, have been used to better understand such systems (Geard and Bullock, 2008; Gross and Sayama, 2009; Sayama et al., 2013).

For instance, Sayama (2020b) introduced a model of opinion dynamics on an adaptive social network in which all agents were subject to some degree of social *conformity* (adjusting opinions in the direction of the local social norm), *homophily* (strengthening connections to agents with similar opinions) and *neophily* (strengthening connections to agents whose opinions are novel with respect to the local social norm). The strength of each of these factors was varied systematically over a series of numerical simulations in order to evaluate their impact on network dynamics.

The study found that to the extent that the population was strongly homophilic, it tended to fragment into a relatively large number of communities with agents from the same community tending to share a similar opinion, but agents from different communities exhibiting divergent opinions that could be extreme relative to the population average[1]. By contrast, a strongly neophilic population tended to form a small number of large communities (often just one), featuring opinions that were moderate and fairly homogeneous (but diverse by comparison with the tightly converged opinions within a single typical homophilic community).

Subsequent work demonstrated that when agents differed in the strength of these social tendencies, i.e., they were *heterogeneous* in terms of their update policies, more extreme and polarised opinions tended to arise (Bullock and Sayama, 2023). This previous work has tended to consider scenarios in which the social network is fully connected, i.e., every agent has the potential to interact directly with every other agent. This raises the following question: how sensitive is this model to the assumption that agent interactions are not spatially constrained?

Here we extend this agent-based model of opinion dynamics on an adaptive social network to explore scenarios in which agent interactions are to some extent constrained by spatial proximity. We conduct a series of numerical simulations and perform regression analyses to elucidate the effects of spatial network structure on a population's tendency towards extremist communities.

Model

Network Initialisation

Following Sayama (2020b) and Bullock and Sayama (2023), the opinions and social structure of a population of agents is represented by a graph, G, comprising a set of nodes, V (where $|V| = N$), connected by a set of weighted, directed edges, E. Each node $i \in V$ represents an individual agent with an opinion $x_i \in \mathbb{R}$. Each node is influenced by each of its network neighbours. The strength of the influence exerted by node j on node i is represented by the weight of network edge $w_{ij} \in \mathbb{R}_{\geq 0}$. Whereas previous studies considered fully connected networks where $|E| = N(N-1)$, here edges are assigned according to one of three spatial recipes:

Ring Lattices: Nodes are distributed uniformly around a ring. Each node is connected directly to the k nodes nearest to it on the ring (excluding itself), assigning a total of $|E| = Nk$ edges. Ring lattices have strong clustering, $c \approx 0.75$, and relatively long characteristic path length, $L \approx \frac{N}{2k}$, but weak community structure due to their perfect uniformity.

[1]In this paper we will tend to describe this kind of between-community polarisation in terms of increased *extremism* which carries negative connotations. However, in some contexts it could equally well be described more positively in terms of increased *diversity* or *innovation*.

Connected Caveman Graphs: Each network node is assigned a "cave" such that every cave has k nodes and caves are distributed equally around a ring. All nodes within the same cave are connected to each other to form a clique. Then K nodes within each cave have one of their edges rewired to a unique node in the clockwise adjacent cave. Like lattices, connected caveman graphs have high clustering ($c = 1$ where $K = 0$) and relatively long path length, but they also have strong community structure, with each cave being readily identifiable as a strong community if K is low.

Random Geometric Graphs: Every node is assigned a location chosen uniformly at random within the unit square. All pairs of nodes separated by a distance less than some threshold d are directly connected, where d is chosen such that the total number of edges, $|E| = Nk$, to obtain comparability with the two recipes above. For the purposes of calculating the distance between two nodes, the embedding space is treated as having toroidal boundary conditions.

For each of these recipes, a degree of random rewiring can also be imposed to erode the influence of spatial constraints on the network topology to some extent. Here, we employ the "conservative" rewiring scheme introduced by Iotti et al. (2017) which conserves the degree distribution of the original network and does not introduce spurious spatial anti-correlations that arise if a more naive degree-conserving random rewiring scheme is employed:

1. For each pair of connected nodes, i, j: With probability p, remove edges ij and ji and add i and j to a rewire list. A node may be added to this list more than once if more than one of its edges is to be rewired. This step is performed for all edges of the initial network, populating the entire rewire list, before proceeding to step 2.

2. Until the rewire list is empty: Remove a randomly chosen pair of values, i and j, from the rewire list, where $i \neq j$ and the nodes are not already connected. Add edge ij and ji to the network. Note: This process may fail if the list contains no legal pairs of nodes. In this case the whole rewiring scheme is restarted.

Network Dynamics

Over time, each node's opinion has a tendency to shift towards the weighted average opinion of its local social neighborhood (social conformity) and also to drift at random (noise). Each node's incoming edge weights also change over time such that they tend to reflect the extent to which the node and its upstream neighbour share a similar opinion (homophily), and also tend to reflect the extent to which the upstream neighbour's opinion is distinct from the weighted average opinion of the node's social community (neophily). That is, opinions and edge weights co-evolve over time through four dynamic mechanisms: (1) social conformity, (2) noise, (3) homophily and (4) neophily.

These dynamics are determined as follows:

$$\frac{dx_i}{dt} = c_i\left(\langle x\rangle_i - x_i\right) + \epsilon \tag{1}$$

$$\frac{dw_{ij}}{dt} = h_i F_h(x_i, x_j) + a_i F_a(\langle x\rangle_i, x_j) \tag{2}$$

$$\langle x\rangle_i = \frac{\sum_{j\in N_i} w_{ij} x_j}{\sum_{j\in N_i} w_{ij}} \tag{3}$$

Here, N_i is the set of in-neighbors of node i (i.e., all nodes connected to i); $\langle x\rangle_i$ is the local weighted average opinion, or social norm, perceived by node i; ϵ represents a stochastic fluctuation term that influences node opinions; and c_i, h_i, and a_i are node-specific parameters that determine, respectively, the strength of social conformity, homophily, and neophily specific to node i. Following Bullock and Sayama (2023), node-specific values for c, a and h were each drawn from a parameter-specific random distribution during initialisation and were kept fixed for the duration of each network simulation. In combination, the c_i, h_i and a_i values of a specific node are referred to as its *update policy* and the population mean parameter values, μ_c, μ_h, and μ_a define the network's *mean update policy*.

Behavioural functions F_h and F_a determine the rate of edge weight change based on opinion distance, defined as follows:

$$F_h(x_i, x_j) = \theta_h - |x_i - x_j| \tag{4}$$

$$F_a(\langle x\rangle_i, x_j) = |\langle x\rangle_i - x_j| - \theta_a \tag{5}$$

Here θ_h and θ_a are fixed population-wide parameters that act as threshold opinion distances separating a regime in which weights from an upstream neighbour are strengthened from a regime in which they are weakened.

Note that w_{ij} is bounded to be non-negative, i.e., any negative values are rounded up to zero. Where a node's incoming edge weights are all zero, its local weighted average opinion is undefined and neither the influence of conformity nor neophily are applicable. Note also that a pair of nodes are neighbours if and only if they were connected by an edge during network construction.

The above adaptive social network model was implemented in Python 3.12 with the NetworkX package.[2]

Simulations

Simulation Settings

For all simulation runs: $N = 1000$, $\theta_h = 0.03$, and $\theta_a = 0.03$. During network initialisation every connected pair of nodes was assigned a directed edge in each direction with a weight randomly sampled from the uniform distribution $[0, 1]$, and each node was assigned a random opinion sampled from the standard normal distribution $\mathcal{N}(0, 1^2)$.

[2]The simulator code is available from the author upon request.

Node-specific parameter values were assigned using one of two methods. *Heterogeneous* agent populations were initialised such that for each node, i, values of c_i, h_i and a_i were randomly sampled from the uniform distribution $[0.01, 0.3]$, i.e., the original range of parameter values explored by Sayama (2020b). For simulations exploring *Homogeneous* equivalents, each c_i, h_i and a_i value was set to be the mean of the Heterogeneous network to which the Homogeneous network was being compared.

Each simulation instantiation employed a simple Euler forward integration method with time interval $\Delta t = 0.1$ for t running from 0 to 100. The stochastic effect of ϵ was simulated by adding a random number sampled from $\mathcal{N}(0, 0.1^2)$ to each x_i at every interval Δt.

Community Structure and Extremism

The Louvain modularity maximization method (Blondel et al., 2008) was employed to assign each node in the final network configuration to one of a set of non-overlapping communities. This method requires that the network be undirected. Consequently, after each simulation run was completed, an undirected network G' was constructed to be the equivalent of the final state of network G at $t = 100$. The weight of each undirected edge in G' was set to be equal to the mean weight of the directed edges between the same pair of nodes in G, i.e., $w_{ij}^{G'} = \frac{1}{2}(w_{ij}^{G} + w_{ji}^{G})$. The community structure of G' was then determined using the Louvain modularity maximization method and the resultant assignment of nodes to communities was applied to G for the purposes of all subsequent analyses.

After determining the mean opinion in each of the communities identified in the final network configuration, the degree of extremism exhibited was measured by calculating the *community mean opinion range*, i.e., the absolute difference between the mean opinion of the agents in the community with the highest mean opinion and the mean opinion of the agents in the community with the lowest mean opinion.

Results

Uniform Spatial Networks

First, we compare the degree of extremism exhibited in agent communities embedded on two canonical classes of spatial graph that lack strong community structure: *ring lattices* (which are constructed over a strictly uniform distribution of nodes around a ring) and *random geometric graphs* (which are constructed over a uniform *random* distribution of nodes on the unit square). In each case, we explore how eroding or relaxing the spatial constraints on the network impacts the degree of extremism in the final network communities. We implement this erosion by imposing a degree of random rewiring controlled by parameter p. Where $p = 0$, no rewiring occurs and network topology is fully determined by the original spatial constraints, ensuring that the network has a high clustering coefficient and a long characteristic path

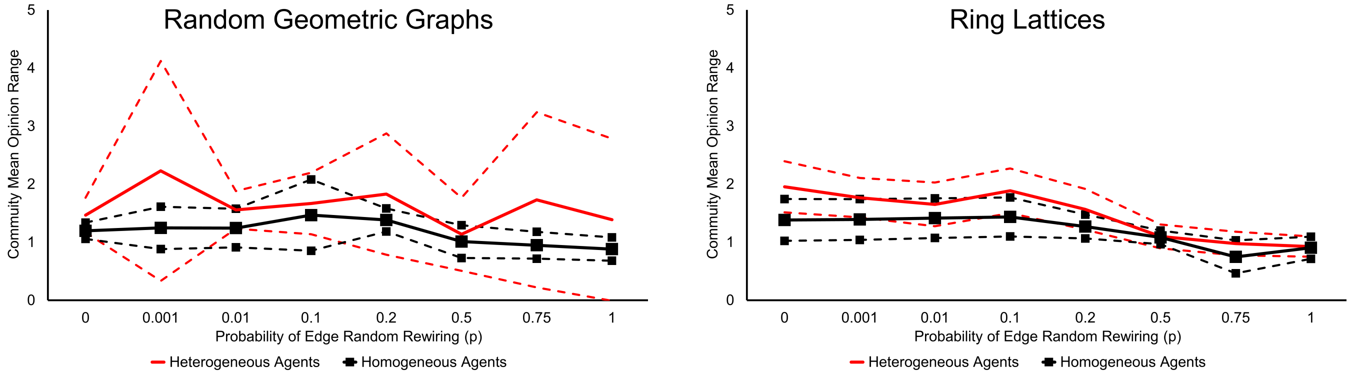

Figure 1: The influence of spatial constraint erosion (due to random rewiring with probability, p) on the degree of extremism (measured in terms of community mean opinion range) for $N = 1000$ agents embedded on *(left)* ten random geometric graphs or *(right)* ten ring lattices with degree $k = 20$. Dashed lines indicate standard deviations around the mean.

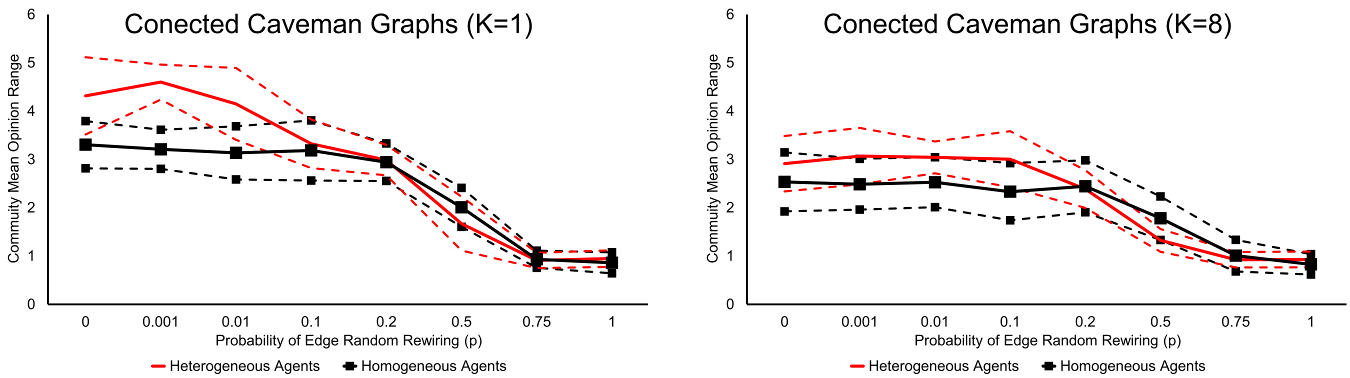

Figure 2: The influence of spatial constraint erosion (due to random rewiring with probability, p) on the degree of extremism (measured in terms of community mean opinion range) for $N = 1000$ agents embedded on ten connected caveman graphs featuring 20-node caves with either low inter-cave connectivity (*left*, $K = 1$) or stronger inter-cave connectivity (*right*, $K = 8$).

length. Conversely, where $p = 1$ all of the network's edges are replaced by randomly assigned edges, and the influence of space is entirely extinguished, resulting in a network with a very low clustering coefficient and a much lower characteristic path length (but exactly the same degree distribution as the original network).

Figure 1 shows that Heterogeneous agent networks tend to organise into somewhat more extreme communities than equivalent Homogeneous networks regardless of the degree to which the spatiality of the network structure has been eroded and that this is true for both ring lattices ($t = 5.08$, df=79, $p < 10^{-3}$) and random geometric graphs ($t = 3.41$, df=79, $p < 10^{-3}$). For the latter, the random distribution of nodes leads to higher variability in the degree of extremism, particularly where the agents were also heterogeneous in their update policy.

There is a mild but significant negative relationship between p, the degree to which spatial constraints are eroded, and the extremism of the final network for homogeneous random geometric graphs ($r^2 = 0.18$, df=78, $p < 10^{-4}$) and for both homogeneous ring lattices ($r^2 = 0.37$, df=78, $p < 10^{-8}$) and heterogeneous ring lattices ($r^2 = 0.524$, df=78, $p < 10^{-13}$). These results indicate that stronger spatial constraints have a tendency to *encourage* more extreme outcomes. However, the size of the effect is not large and the variability of the heterogeneous random geometric graph outcomes prevents any relationship from being found in that case.

Spatial Networks with Community Structure

The impact of spatial structure on the degree of extremism and polarisation is more evident for connected caveman graphs, which feature a larger number of stronger spatial communities than both ring lattices and random geometric graphs. Figure 2 depicts results for 1000-node connected caveman graphs with weakly connected 20-node caves ($K = 1$) and more strongly connected 20-node caves ($K = 8$). In each case there is a clear fall in extremism as the spatial constraints on the networks are relaxed. This is true for a range of inter-cave connectivity levels (see Figure 3). In each case, the most extreme networks are those featur-

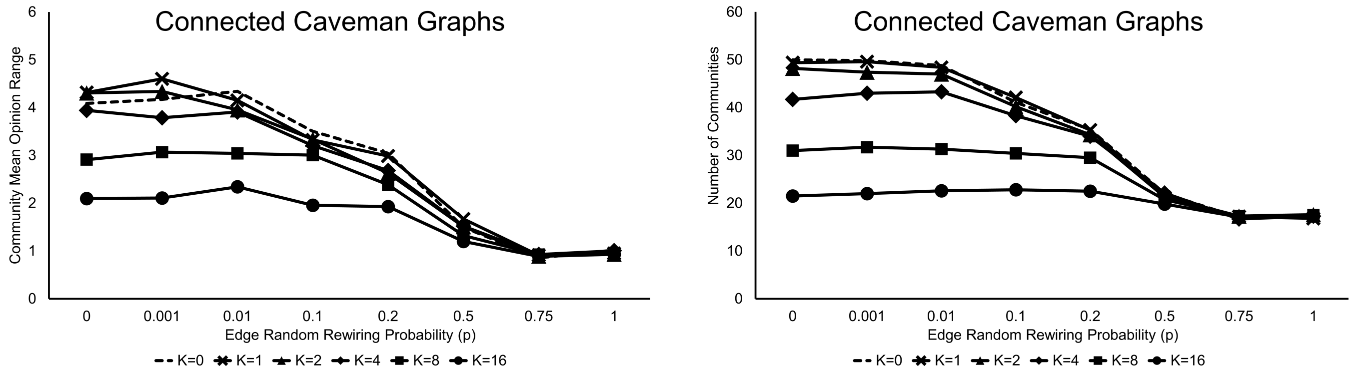

Figure 3: The influence of spatial constraint erosion (due to random rewiring with probability, p) on *(left)* the degree of extremism (measured in terms of community mean opinion range) and *(right)* the number of communities identified in the final network for $N = 1000$ heterogeneous agents embedded on ten connected caveman graphs featuring caves of size $k = 20$ and differing degrees of inter-cave connectivity, $K \in \{0, 1, 2, 4, 8, 16\}$.

Table 1: Multiple linear regression of a connected caveman network's *community mean opinion range* (CMOR) outcome on *(left)* its structural parameters (K, p) and their interaction (bottom row) and *(right)* the number (C) and strength (m) of its identified communities and their interaction (bottom row), for both Homogeneous agent networks and Heterogeneous agent networks. Each model was built on the data plotted in Figure 3. Statistically significant coefficients are indicated with asterisks (*: $p < 10^{-2}$; **: $p < 10^{-3}$; ***: $p < 10^{-4}$; etc.).

Network	Homogeneous	Heterogeneous
Outcome	CMOR	CMOR
const.	3.484^{***}	4.090^{***}
K	-0.104^{***}	-0.124^{***}
p	-2.863^{***}	-3.780^{***}
$K \times p$	0.124^{***}	0.149^{***}

Network	Homogeneous	Heterogeneous
Outcome	CMOR	CMOR
const.	-0.303	-0.555
C	0.070^{***}	0.054^{*}
m	1.065^{***}	1.376^{***}
$C \times m$	-0.020	0.018

ing agents with heterogeneous update policies embedded on connected caveman graphs that have only been subjected to low levels of random rewiring (low p).

Increasing p has several important effects on connected caveman network structure. First, random rewiring reduces the network's average clustering coefficient (the proportion of pairs of nodes with a common neighbour that are themselves neighbours) and its characteristic path length (the average length of the shortest paths between every pair of nodes). The latter tends to decrease more rapidly than the former, resulting in a "small world effect" for low positive values of p where the clustering coefficient remains high relative to the original un-rewired network, but the characteristic path length is much reduced. There is some evidence of a subtle small-world effect on extremism to be found in the observation that the highest values for community mean opinion range tend to occur for $0 < p \ll 1$, but this effect is not statistically significant. Second, random rewiring tends to blur the boundaries between caves and consequently lowers the number of distinct communities, C, in the best partition identified by the Louvain community detection method, and lowers the modularity score, m, of this partition, i.e., the community structure is weakened by random rewiring. Simple regression analyses were employed to identify which of these structural consequences of random rewiring are responsible for the drop in extremism.

The influence of model parameters and network structure on the level of extremism exhibited was explored using multiple linear regression models (see Table 1). Regression of a network's extremism outcome measure on its network parameters (K and p) confirms that extremism is significantly related to both factors and that they also exhibit a significant amplificatory interaction effect. Moreover, a second multiple linear regression model demonstrates that the number of communities identified within a network (C) and the strength of these communities (m) are both significant positive predictors of the final level of extremism. When the average clustering coefficient and characteristic path length are made available as additional factors within the multiple linear regression model they are not identified as significant predictors of extremism. Consequently, we conclude that manipulating K and p during network construction has an impact of the community structure of the resultant network (as reflected in C and m) and that the strength of this community structure is what tends to influence the extremism of the final network configuration. The striking relationship between the fall in the number of communities identified in heterogeneous connected caveman networks and the attendant fall in extremism is depicted in Figure 3.

Discussion

Imposing the "death of distance" (Cairncross, 2001) on an adaptive spatial social network model does not result in an increase in polarisation and extremism within the network's communities. Conversely, the introduction of random (typically long-distance) connections into spatially constrained networks has tended to erode the integrity of spatial communities and the boundaries between them, thereby encouraging more *moderate* opinions to predominate. The introduction of indiscriminate connectivity does not tend to allow extreme individuals to connect to each other and form extreme communities. Rather, it is the presence of strong spatial community structure that allows some of the relatively isolated fragments of the population to settle on extreme opinions and resist conforming to the wider population-level social norm.

Three aspects of the model could be improved in order to better explore the influence of online social networks on opinion dynamics. Each is the intended focus of future work.

First, while the model as it stands allows structural change in the network to take place as a result of weight changes due to homophily and neophily, it is the case that two nodes can only directly influence one another if they happen to be connected together during network initialisation. The model's network dynamics do not allow nodes to shed connections to existing neighbours or gain connections to new neighbours, despite this being a key feature of real world social network evolution.

Second, during the rewiring process, the edges removed and the edges introduced are selected at random. One key aspect of online social networks is their ability to allow like-minded individuals to connect with one another. A rewiring scheme in which randomly chosen edges that tend to link nodes with dissimilar opinions tend to be replaced by edges connecting randomly chosen nodes with similar opinions would allow this aspect of real-world social network evolution to be captured.

Third, while network agents vary in the *rate* at which they alter their opinion to conform to the social norm around them (controlled by their c_i parameter) all agents attempt to conform to a maximal *extent*. That is, every agent will reduce the difference between its own opinion and that of the social norm around them gradually until this difference is zero. In reality, some agents may be comfortable with an opinion that deviates from the norm to some degree. Moreover, some agents may actively seek to achieve and maintain an opinion that deviates from the social norm in order to be interesting, distinctive or controversial. Introducing variation in the *extent* to which agents conform (as well as the rate of conformity) could enable exploration of the impact of "influencer" agents on social network dynamics.[3]

Conclusions

In this paper, we have extended an existing computational agent-based model of adaptive social network dynamics by introducing spatial constraints on agent interactions and investigated the impact of this spatial structure on the way in which extreme communities can arise through a combination of social conformity, homophily and neophily. Our results demonstrate that spatial organisation can have a systematic influence on the outcome of social network formation, particularly when it introduces strong community structure within the network and particularly when agent update policies are heterogeneous. Relative to comparable non-spatial networks, networks with strong spatial community structure tend to support communities with a more diverse set of opinions and to allow individuals with extreme opinions to have more influence on their local community.

Acknowledgments

This work was supported by UKRI EPSRC Grant No. EP/Y028392/1: *AI for Collective Intelligence (AI4CI).*

References

Axelrod, R., Daymude, J. J., and Forrest, S. (2021). Preventing extreme polarization of political attitudes. *Proceedings of the National Academy of Sciences*, 118(50):e2102139118.

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008.

Bullock, S. (1998). A continuous evolutionary simulation model of the attainability of honest signalling equilibria. In Adami, C., Belew, R., Kitano, H., and Taylor, C., editors, *Artificial Life VI: Proceedings of the Sixth International Conference on the Synthesis and Simulation of Living Systems*, pages 339–348. MIT Press, Cambridge, MA.

Bullock, S. (2016). "Shit happens": The spontaneous self-organisation of communal boundary latrines via stigmergy in a null model of the European badger, *Meles meles*. In Gershenson, C., Froese, T., Siqueiros, J. M., Aguilar, W., Izquierdo, E. J., and Sayama, H., editors, *Artificial Life XV: Proceedings of The Fifteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 518–525. MIT Press.

[3]Thanks to Jadesola Bejide for discussions on this idea.

Bullock, S. and Buckley, C. L. (2009). Embracing the tyranny of distance: Space as an enabling constraint. *Technoetic Arts*, 7(2):141–152.

Bullock, S. and Cliff, D. (1997). The role of 'hidden preferences' in the artificial co-evolution of symmetrical signals. *Proceedings of the Royal Society of London, Series B*, 264:505–511.

Bullock, S. and Geard, N. (2010). Spatial embedding as an enabling constraint: Introduction to a special issue of complexity on the topic of 'spatial organisation'. *Complexity*, 16(2):8–10.

Bullock, S. and Sayama, H. (2023). Agent heterogeneity mediates extremism in an adaptive social network model. In Iizuka, H., Suzuki, K., Uno, R., Damiano, L., Spychalav, N., Aguilera, M., Izquierdo, E., Suzuki, R., and Baltieri, M., editors, *Proceedings of the Artificial Life Conference 2023 (ALIFE 2023)*, pages 258–266. MIT Press.

Cairncross, F. (1997/2001). *The Death of Distance: How the Communications Revolution is Changing our Lives*. Harvard Business School Press.

Geard, N. and Bullock, S. (2008). Group formation and social evolution: A computational model. In Bullock, S., Noble, J., Watson, R., and Bedau, M., editors, *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, pages 197–203. MIT Press.

Gross, T. and Sayama, H. (2009). *Adaptive Networks*. Springer.

Iotti, B., Antonioni, A., Bullock, S., Darabos, C., Tomassini, M., and Giacobini, M. (2017). Infection dynamics on spatial small-world network models. *Physical Review E*, 95(5–1):052316.

Jacyno, M., Bullock, S., Geard, N., Payne, T. R., and Luck, M. (2013). Self-organising agent communities for autonomic resource management. *Adaptive Behavior*, 21(1):3–28.

Jacyno, M., Bullock, S., Luck, M., and Payne, T. R. (2009). Emergent service provisioning and demand estimation through self-organizing agent communities. In Sierra, C., Castelfranchi, C., Decker, K. S., and Sichman, J. S., editors, *Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, pages 481–488. ACM.

Levin, S. A., Milner, H. V., and Perrings, C. (2021). The dynamics of political polarization. *Proceedings of the National Academy of Sciences*, 118(50):e2116950118.

Pitonakova, L., Crowder, R., and Bullock, S. (2016a). Information flow principles for plasticity in foraging robot swarms. *Swarm Intelligence*, 10(1):33–63.

Pitonakova, L., Crowder, R., and Bullock, S. (2016b). Task allocation in foraging robot swarms: The role of information sharing. In Gershenson, C., Froese, T., Siqueiros, J. M., Aguilar, W., Izquierdo, E. J., and Sayama, H., editors, *Artificial Life XV: Proceedings of The Fifteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 306–313. MIT Press.

Pitonakova, L., Crowder, R., and Bullock, S. (2018). The Information-Cost-Reward framework for understanding robot swarm foraging. *Swarm Intelligence*, 12(1):71–96.

Santos, F. P., Lelkes, Y., and Levin, S. A. (2021). Link recommendation algorithms and dynamics of polarization in online social networks. *Proceedings of the National Academy of Sciences*, 118(50):e2102141118.

Sayama, H. (2020a). Enhanced ability of information gathering may intensify disagreement among groups. *Physical Review E*, 102:012303.

Sayama, H. (2020b). Extreme ideas emerging from social conformity and homophily: An adaptive social network model. In Bongard, J., Lovato, J., Hebert-Dufrésne, L., Dasari, R., and Soros, L., editors, *Proceedings of the 2020 International Conference on Artificial Life*, pages 113–120. MIT Press.

Sayama, H., Pestov, I., Schmidt, J., Bush, B. J., Wong, C., Yamanoi, J., and Gross, T. (2013). Modeling complex systems with adaptive networks. *Computers & Mathematics with Applications*, 65(10):1645–1664.

zu Erbach-Schoenberg, E., Bullock, S., and Brailsford, S. (2013). A model of spatially constrained social network dynamics. *Social Science Computer Review*, 32(3):373–392.