

The impact of peer pressure: extending Axelrod’s model on cultural polarisation

Rhodri L. Morris, Liam D. Turner, Roger M. Whitaker
School of Computer Science & Informatics,
Cardiff University
Cardiff, UK
{MorrisRL6, TurnerL9, WhitakerRM}@cardiff.ac.uk

Cheryl Giammanco
U.S. Army Combat Capabilities Development Command,
Army Research Laboratory
Aberdeen Proving Ground, MD, USA
cheryl.a.giammanco.civ@mail.mil

Abstract—Culture represents the broad range of things over which people influence each other, and frequently contributes to the behaviour, interaction and outlook of groups. Although it has been studied in the context of humans, it is also relevant to future intelligent cognitive systems, that could have the capability to update their disposition and strategy based on the influence of others.

In this work we transfer concepts from social sciences to the computing sciences and examine the effect of peer influence on culture. We consider the notion of “peer pressure”, being the combined effect from all an individual’s neighbours exerting influence at the same time, and also through influence flowing from indirect sources. This approach is derived using Social Impact Theory. We benchmark this against the cultural polarisation model from Axelrod, which involves influence being restricted to dyadic interactions between agents.

We find that peer pressure provides complex contagion with a significant impact on cultural evolution. Greater cultural diversity is maintained, with indirect paths mitigating this by effectively forming disruptive weak links. This reaffirms that maintaining diversity in social ties, as well as a wide breadth, supports the mitigation of cultural isolation and polarisation. The model provides a platform to explore culture in a wide range of further scenarios, including electronic, coalition and organisational contexts.

Keywords-Computational Modelling, Cultural polarisation, Agent-based Modelling, Homophily

I. INTRODUCTION

Humans influence each other across a wide spectrum of issues and activity, from ideas, to politics, fashion, behaviour and social norms. Broadly speaking, we can refer to the things over which people influence each other as *culture* [1]. This is enabled by an underlying social network, through which individuals interact and may adopt alternative cultural elements or reaffirm existing ones [2]. As part of this, influence can spread across the social network [3], enabling a degree of indirect influence between nodes.

Across human evolution, this is a never-ending phenomenon that contributes to diverse social issues such as the formation of groups, polarisation, viral effects on social media and the spread of ideas. However, with the acceleration of artificial intelligence, one can also envision a future where devices or agents are equipped not just to learn from first-hand data acquisition, but are able to navigate their social world by potentially learning from others, including updating

meta-level issues such as their behavioural dispositions and their view of the cyber-world in which they function (e.g., threat levels).

With this in mind, in this paper we are interested in further examining the transfer of cultural concepts from social sciences to the computing sciences [4]. Specifically, we computationally model cultural evolution over a network structure, while taking into account the effect of social influence from neighbours. Our work addresses two main issues. Firstly, we address the notion of “peer pressure”, being the combined effect from all an individual’s neighbours exerting influence at the same time. Secondly, we address indirect influence that an individual C may exert on A , due to C ’s relationship with B , who holds a direct relationship with A . Here for example, A might have a vague knowledge of C ’s presence, but the impact of C on A is mediated via B .

The theoretical basis for our model comes from Social Impact Theory [5], a theory that has been adopted to emulate the larger scale effects of a range of different social and psychological processes (e.g., [6], [7], [8], [9]). We apply this as part of a computational model to consider the combined effect that one’s neighbours have in the social influence of personal culture. To benchmark our work we adopt a model underlying the seminal work of Robert Axelrod [1], who used computational modelling to understand why populations maintain cultural diversity, rather than always converging to a homogeneous state.

Overall, our paper presents two key contributions. Firstly, we find that our model allows for a compound effect to be observed when neighbouring views are collectively considered, in comparison to pairwise dyadic interactions in Axelrod’s model (described in Section II-A). Secondly, we determine the relative effects of alternative social influences (described in Section II-B and III) to which individuals are exposed and find that direct and indirect social ties can mitigate cultural polarisation.

II. RELATED WORK

A. Axelrod’s approach to cultural evolution

The approach adopted by Axelrod [1], and in many subsequent related papers, involves describing an individual

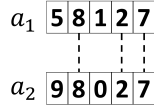


Figure 1. An example of two agents (a_1 and a_2) with their associated feature vectors. We assume $F = 5$ and $q = 10$, using traits denoted 0-9. These agents have three common features (positions 2, 4 and 5). Under Axelrod’s model, a_1 may copy a dissimilar trait from a_2 with probability $p = 0.6$.

agent’s culture as a vector of *features*, where each feature has a value from a set of discrete *traits* unique to that feature. Each agent has F cultural features ($\sigma_1, \sigma_2, \dots, \sigma_F$), and each feature (which could represent, for example, a different taste or behaviour) takes its value from a range of q possible traits, where $\sigma_f = 1, \dots, q$ for feature f .

The network structure involves agents being present in a bounded regular lattice configuration, allowing interactions with their north/south/east/west neighbours (i.e., Von Neumann neighbourhood) unless the agent has a boundary position. Axelrod’s model functions by a random agent A being selected to choose a random neighbour B to interact with. A then decides whether to copy a dissimilar feature from B . This is governed by the overall similarity of A and B , which can be expressed as the proportion of features where A and B hold the same trait, denoted p . Then A chooses a dissimilar feature to copy from B with probability p . An example is shown in Figure 1. Thus when two neighbouring agents have more features in common, they are more likely to copy an additional trait from each other, thereby reinforcing the role that homophily [10], [11], [12] plays in attracting those that are similar and supporting their assimilation.

This elegant model allows cultural regions to form, representing connected sub-networks where all agents have identical features. Interestingly, polarisation is inherent in the model: changes cannot take place when an agent is either completely similar or dissimilar to its neighbours. Thus cultural regions become stable when they have no features in common with their bordering cultural regions. Axelrod [1] determines the effect of the number of features and traits on the presence of cultural regions.

There have been numerous extensions to this fundamental model of cultural evolution. For example, Castellano et al [13] define a critical threshold where cultural diversity is more common than global monoculture. Klemm et al [14] introduce ‘cultural drift’ by the inclusion of random perturbations. Allowing agents to move within the regular lattice, based on similarity, has also been considered [15], resonating with the work of Schelling [16]. Centola et al [17] also consider changing the underlying structure of the model by changing links when there are no common features between agents, thereby observing co-evolution of network structure. Further consideration has also been given to the influence of features [18] that particular agents hold. For example two agents may

be able to influence each other on sport but not politics. The collective influence of neighbours has received relatively little attention, despite it being of significant importance in social interactions. Flache and Macy [19], Parisi et al [20] each use models which incorporate some form of compound influence, albeit without the effect of influence diminishing with distance as described in Section II-B.

B. Social Impact Theory

Latané’s fundamental work of ‘The Psychology of Social Impact’ [5] proposes a theory of how an individual influences, and is influenced by, their peers. This characterisation involves three variables: the number of individuals providing influence (denoted N); their strength (denoted S) reflecting the notion of power through factors such as status, resources or in-group membership; their immediacy (denoted I), broadly relating to the proximity in temporal, physical or social space. In its general form, the magnitude of social impact from peer influence is a function $f(SIN)$.

This general approach has been tailored to consider a range of settings, including influence on language change [6], culture based on specific issues [21] and attitudes [8]. In our model (discussed in Section III), inspired by the proposal by Nowak et al [22] in modelling opinion formation, we develop a multi-dimensional model of culture. This work developed a cellular automaton modelling the change of binary opinions using forces of support and persuasion from nearby actors to influence whether an agent will retain or change their opinion respectively, with influence on the target decreasing as distance increases.

III. MODEL

Let N denote the set of agents, assuming $|N| = n$. We assume that each agent j holds a feature vector $(\sigma_1^j, \sigma_2^j, \dots, \sigma_F^j)$ where there are F features and each feature is populated by one of q traits. Let T denote the set of traits.

The similarity, sim_{ij} between neighbours i and j is defined as follows. Let $s_{ij}^k = 1$ if and only if $\sigma_k^i = \sigma_k^j$, else $s_{ij}^k = 0$. In other words, $s_{ij}^k = 1$ when the k^{th} feature of i and j are the same. Then:

$$sim_{ij} = \sum_{k=1}^F \frac{s_{ij}^k}{F}$$

We weight the edge between i and j , denoted w_{ij} , based on the similarity of i and j , by setting $w_{ij} = sim_{ij}$.

Agents are called at random to consider the social influence that they are subject to from their neighbours, and potentially their neighbour’s neighbours (within an edge distance of l). Therefore we consider the potential paths of length l from the agent i called to update their features based on social influence. We define P_i^l as the set of all paths of length l that end at i (excluding any cycles but allowing multiple paths to a node). Let $p^l \in P_i^l$. Then the weight of path p^l , denoted $w(p^l)$, is the product of all edge weights on p^l . Indirect

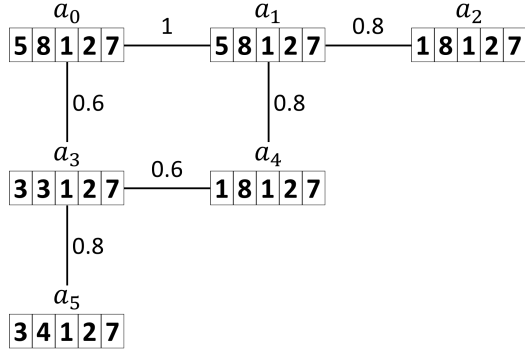


Figure 2. An example of peer-pressure influence using the trait score metric. a_0, \dots, a_5 represent agents with their associated feature vectors ($F = 5, q = 10$). The edge weights denote the similarity metric sim_{ij} . We consider the influence upon the first feature. We assume that the maximum path length for influence $max\ l$ is 2. Three traits provide influence: trait 1 (via a_2 and a_4); trait 3 (via a_3 and a_5) and trait 5 (via a_1). For demonstration, using the trait score metric $ts_{\alpha,k,i}$, we calculate influence from trait 1 (i.e., $\alpha = 1$) upon feature 1 (i.e., $k = 1$) of a_0 (i.e., $i = a_0$). There are three paths of at most length 2 through which a_2 and a_4 invoke influence: namely $a_0 \rightarrow a_1 \rightarrow a_2, a_0 \rightarrow a_1 \rightarrow a_4, a_0 \rightarrow a_3 \rightarrow a_4$. The products of the weights of these paths are 0.8, 0.8 and 0.36 respectively, giving a trait score $ts_{1,1,a_0} = 1.96$. The trait score for the influence of trait 3 on feature 1 of a_0 is $ts_{3,1,a_0} = 1.08$. The trait score for the influence of trait 5 on feature 1 of a_0 is $ts_{5,1,a_0} = 1$. Thus trait 1 is providing the biggest influence in this case.

influence, that is, influence exerted on an agent through the medium of another [23] has been observed in social networks [24], [25]. We use the product of edge weights to mimic the degradation that occurs when influence is mediated by a third party. Each trait α that an agent holds as its k^{th} feature may convey influence on node i through a path p^l . The more similar the agents are along the path, the greater the scope for influence on agent i . We assess this using a *trait score* metric, denoted $ts_{\alpha,k,i}$ where

$$ts_{\alpha,k,i} = \sum_{l=1}^{max\ l} \sum_{p^l \in P_i^l} w(p^l) \delta_{k,\alpha}(p^l)$$

such that $\delta_{k,\alpha}(p^l)$ is a binary variable that is 1 if and only if the source node of the path p^l has α as its k^{th} feature (i.e., if and only if $\sigma_k^j = \alpha$ where j is the source node for the path).

For each feature k , where $k = 1, \dots, F$ the selected agent calculates the trait score for each possible trait in T and selects the value of $\alpha' \in T$ that gives the maximal result. In other words, $\alpha' \in T$ is determined such that $ts_{\alpha',k,i} \geq ts_{\alpha,k,i}, \forall \alpha \in T$. Then σ_k^i is set to α' . This is carried out for all features k , with the target agent i responding to the strongest trait α influencing it. An example of this process is given in Figure 2.

Through the trait score metric, the strength of a relationship between two agents is modelled by their similarity, as represented by edge weight. The immediacy of impact is governed by the maximum path length ($max\ l$) over which

social influences are assumed to be conveyed, and the number of concurrent influences is represented by influences from all the agents present in paths of length l .

Paths remain fixed in the model and as similarity between a pair of agents changes, the influence on all associated paths can be updated at that point. This means that computational complexity need not pose a significant issue for the model.

IV. RESULTS

We place the agents in a 10×10 regular lattice. While unlikely to be realistic for a social network, this structure allows for direct comparison with Axelrod's model of cultural polarisation [1] and subsequent extensions [13], [15]. Agents are created with randomly assigned traits. For each set of parameters discussed below, we perform the simulation until it reaches a stable state where no further changes of traits take place. We examine these end-states over several runs and report on the mean value.

A. Consistency with Axelrod's dyadic-based model

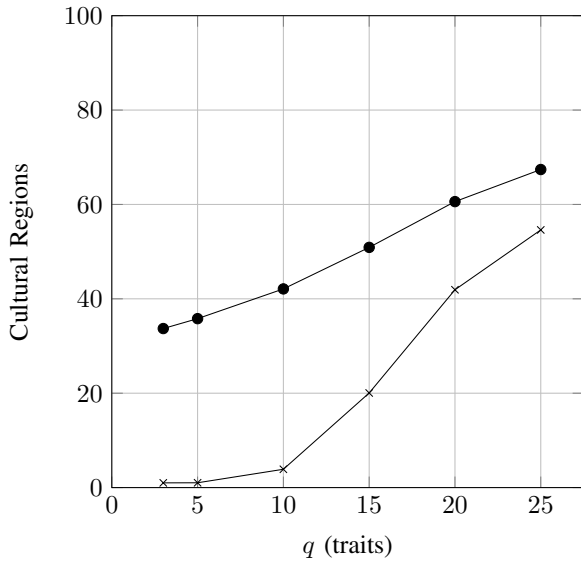
Firstly we verify that our implementation of Axelrod's dyadic-based model [1] has the same characteristics when varying the number of features F or traits q , as also considered in subsequent related models [13], [15], [17]. This is affirmed in Figure 3, where we find that a greater number of features increases the chances of two agents having a feature in common, promoting convergence and resulting in a smaller number of different cultural regions (Figure 3(b)). A greater number of traits for each feature reduces the chances of two agents having the same trait value for a given feature (Figure 3(a)), which is also consistent with previous work.

B. The role of features and traits in cultural diversity

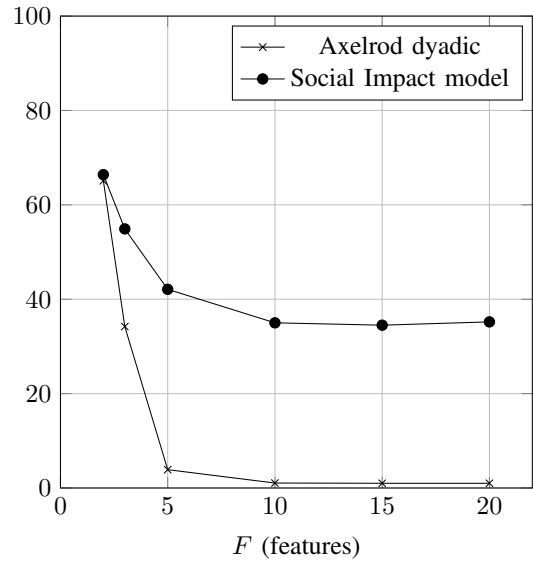
The effects of varying the number of cultural features (F) and the availability of alternative traits q is shown in Figure 3. Our model appears to have a general mediating effect on the number of cultural regions, as compared to Axelrod's dyadic interaction model of cultural influence. As the number of features increases, this provides both models with greater opportunity for similarity to be established, thereby promoting a general reduction in the number of cultural regions (Figure 3b). While this is pronounced for dyadic interactions, the complex contagion present in the social impact inspired model appears to fuel diversity and counters polarisation. Equally, even when the number of traits is small, our model of compound influence is able to sustain a significant number of cultural regions, as compared to dyadic interactions (Figure 3a).

C. The presence of cultural regions and cultural zones

Local convergence and global polarisation are inevitable in Axelrod's dyadic interaction model. However, where multiple sources simultaneously influence an agent, simulations show

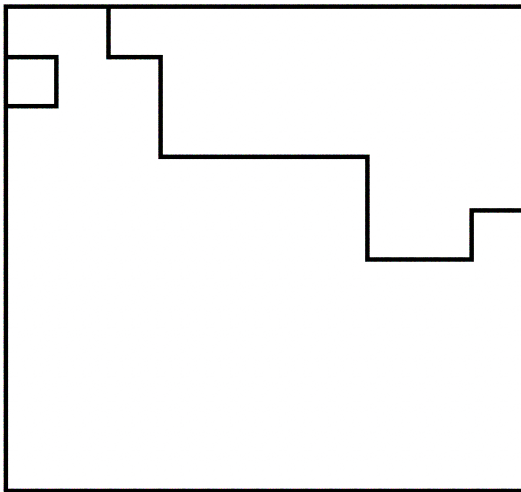


(a) Varying q . $F = 5$, distance = 1

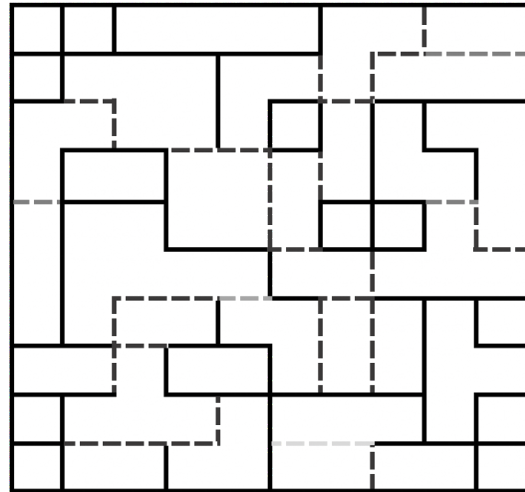


(b) Varying F . $q = 10$, distance = 1

Figure 3. The effect of increasing the number of possible traits q on the number of cultural regions (a), and the effect of increasing the number of features F on the number of cultural regions (b) for both Axelrod's dyadic model of cultural influence and the Social Impact model of peer pressure.



(a) Axelrod's model



(b) Social Impact inspired model

Figure 4. A unit-square representation of agents positioned on a 10×10 regular lattice. Solid lines indicate boundaries between culturally dissimilar agents; lighter lines indicate a greater level of similarity between neighbours. The absence of any dividing line indicates complete cultural similarity. Both (a) and (b) were run with the same starting cultures and random seeds, with 100 agents, 5 features and 10 traits and the distance of influence limited to 1. Axelrod's model (a) stabilises at 3 regions of culturally identical agents, yet each region being completely polarised from the next. However, when agents take into account multiple influences simultaneously, stability can occur with mixed cultural zones still existing (b).

that a stable state may be reached where agents have a mix of traits drawn from bordering regions. An example of this is shown in Figure 4. This means that on average, the number of cultural regions (which Axelrod defines as a set of contiguously sited agents with an identical culture) is therefore usually higher under a model of multiple influences. Here, agents exhibiting a mix of cultures provide additional regions. This is consistent with the findings of

Flache and Macy [19] when implementing other models also incorporating multiple social influences.

In addition to cultural regions, Axelrod defines a *cultural zone* as a set of contiguously located agents where each has a neighbour with at least one feature in common. In Axelrod's model, stability is reached when the number of cultural zones matches the number of cultural regions exactly - no other changes can take place. In our model of multiple influences,

this is not necessarily the case. Cultural zones persist across regions where agents neighbour compatible, but not identical, cultures. This is also evident in Figure 4, where different levels of similarity are evident. Thus a more diverse set of cultural features are maintained.

D. The role of distance in cultural diversity

Interestingly, increasing the distance over which influence may immediately travel (controlled by $max\ l$) causes a reduction in the number of distinct cultural regions, promoting local convergence and global polarisation, as shown in Figure 5. This was also observed in Axelrod’s model [1]. Here, as neighbourhood size increases, those interactions with further agents take place on the same basis as interactions with the closest neighbours. In this model culture may spread by simple contagion; a single interaction with any node within the permitted distance is sufficient for a trait to be copied. However, in our social impact inspired model culture spreads by complex contagion; several influencing agents are often required for a trait to be copied. This may explain the less pronounced drop in cultural regions observed in Figure 5.

Unlike [1] however, our model also incorporates a diminishing effect on influence as distance increases, which represents the mediating effects from indirect communication. While not represented by physical links in the graph, the ability for an agent to be influenced by another several hops away may be thought of conceptually as a form of ‘weak tie’ as originally proposed by Granovetter [26]. In a regular lattice, extending the possible distance of influence will often increase the number of paths accessible to an influencing node. In effect, this widens the ‘bridge’ of weak ties, increasing the possibility of a culture spreading through complex contagion [27].

E. Time to stabilisation

Direct comparison of time-to-stability between the two models is difficult as the typical number of interactions are on very different scales. A single interaction in our model entails an agent examining - and potentially copying - from many neighbours on several features simultaneously. The dyadic model on the other hand involves an agent interacting with only one neighbour per event, and so requires several interactions before it can be said to have been influenced by all neighbours. However, our model does appear to have a more consistent run time due to the stabilising effects of both compound and indirect influence. Across the parameters illustrated in Figure 3, our model typically stabilises at between 550 and 1800 events, with a relative standard deviation of 22%. Axelrod’s model typically stabilises at between 1800 and 225500 events with a relative standard deviation of 72%.

V. DISCUSSION

The results identify that there is a considerable effect on cultural evolution through peer pressure from multiple

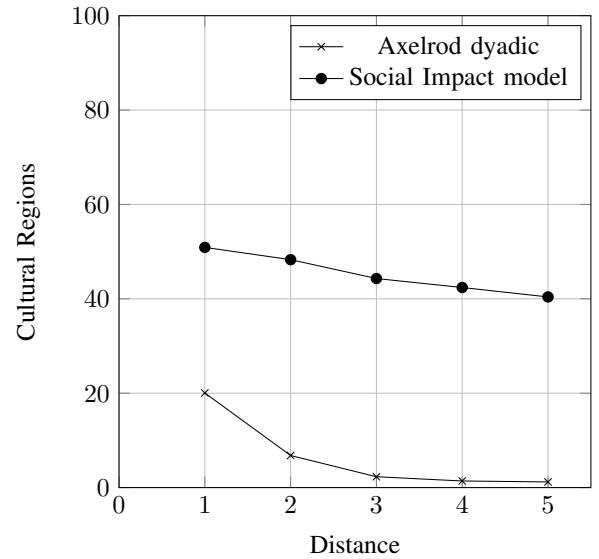


Figure 5. The effect of varying maximum distance over which influence can travel, using $q = 15$, $F = 5$.

simultaneous sources of influence. While we acknowledge that there are likely alternative approaches to modelling the combined effects of peer influence, the simple approach that we have developed, while taking into account Social Impact theory, offers a useful benchmark and basis for future development. In particular, we note that consistent with the model developed by Axelrod, our approach involves the agent’s assimilation with others being mediated by *a-priori* cultural similarity.

We take from the results that restricted cultural influences, such as attained through dyadic interaction, provide a sequence of pairwise interactions where individuals are attracted through similarity and assimilate in a step-wise process that leads to larger groups of identical agents, in terms of their culture. This leads to a generally smaller number of distinct cultural regions that coexist with no overlap. In essence, individuals are limited in sustaining attraction to multiple distinct others through this mechanism.

In contrast, the peer pressure model provides a form of complex contagion, where the assimilation towards a particular feature is a consequence of the influence from a range agents, who may have limited similarity among themselves. This means that when an agent updates a feature through the peer pressure model, it is assimilating with multiple but potentially diverse agents simultaneously. This allows cultural diversity to be maintained: we observe cultural regions emerging within zones (as in the Axelrod model) but these zones do not become homogenous, instead retaining a number of cultural regions with overlapping features. This phenomenon is readily observed in the human social world for example, where a juxtaposition of different cultures leads to

a new culture in its own right. The slower rate of assimilation from multiple sources is consistent with the phenomenon of complex contagion observed by Centola and Macy [27].

A further interesting aspect of the peer pressure model we have introduced is the effect of indirect cultural influences that flow through paths of a given maximum length within the social network (i.e., lattice structure). In some sense, these flows are acting as “weak links”, providing a reinforcement of influence from afar. Because the influence associated with paths is a product of the edge weights along the path, which in turn represents the similarity between the corresponding pairs of agents, only the influence mediated through a path of mutually similar agents retains considerable strength. However, the cumulative affect of many low weight paths, which occur when longer maximum path lengths are invoked, can be significant. Weak links are well known to have important consequences for social and structural connectivity, and our results align with this proposition.

Finally, we believe that the results and formulation of the model may support multiple purposes. On the one hand, the model has potential to provide insights into the human world, where interactions are known to be well-structured by some form of social network. Different organisational networks are such an example; the application of network science to the study of organisational performance is already well established [28], [29]. Alternatively, if the agents are considered to be purely machine-driven entities, there is the potential to understand how different design choices, in how agents may choose to learn from each other, result in potentially different forms of collective cultures, and how these cultures are shared. Examination of these collective behaviours has provided inspiration for heuristic methods [30], [31].

VI. FUTURE WORK

Although a regular lattice is a useful simplification, network topologies more pertinent to social or organisational structures may be considered. For example, the Axelrod model has already been implemented on scale-free [32], small-world [33] and Erdős-Rényi random graphs [34]. In addition, hierarchical structures may allow investigation of how culture can spread through an organisation: the effect of ‘horizontal’ ties between departments or teams, or edge weights being adjusted for higher agent status or reputation. Evolving network structures have been developed for the dyadic Axelrod model [17]; similar extensions may reveal interesting behaviours in a model of compound influence. Depending on how rewiring of edges takes place, changing links may either serve to reinforce existing traits or provide the additional bridge needed for an agent to adopt a new culture.

In both the Axelrod and peer pressure model, an agent’s trait values are discrete; there is no notion of distance between traits. Traits may also be modelled as a continuous sequence,

as in several works on opinion formation [2], [35], [36]. A scale of opinions represented by a sufficiently large features vector may also be used to produce a similar behaviour.

A possible limitation of this work is an under-representation of the importance of an agent’s existing state when considering possible traits. Models of opinion formation often incorporate a force of resistance acting against the probability of change [23], sometimes implemented by means of a similarity threshold which needs to be exceeded if change is to occur [34]. Agent susceptibility to change need not be an homogeneous behaviour either; individuals may display differing levels of resistance to change with committed agents having an effect on the rate of assimilation [34] or even allowing a minority to convert a majority [37].

VII. CONCLUSIONS

In this paper, we extend Axelrod’s Dissemination of Culture model using Latané’s Social Impact Theory [5] as a basis for simulating “peer pressure” on an individual from the combined effect of its neighbours. Concepts from social sciences have been transferred to the computing sciences and we have examined the impact of peer influence on culture. We have established the power of peer pressure through its role as a complex contagion, and contrasted this model against the seminal work of Axelrod [1]. Our findings show that peer pressure invokes a significant increase in the number of stable regions that emerge, while taking into account diminishing influence by social distance. Although simplifying assumptions have been imposed, such as a regular lattice structure and fixed links between agents, the model introduced provides a useful basis for further experimentation, including alternative network structures, heterogeneous agent behaviour and alternative models for influence taking hold upon individuals. The outcomes of such modelling are potentially relevant to organisational behaviour as well as understanding inter-group relations, in terms of coherence and divergence of common views.

ACKNOWLEDGMENT

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

REFERENCES

- [1] R. Axelrod, “The dissemination of culture: A model with local convergence and global polarization,” *Journal of conflict resolution*, vol. 41, no. 2, pp. 203–226, 1997.

- [2] N. E. Friedkin and E. C. Johnsen, "Social influence and opinions," *Journal of Mathematical Sociology*, vol. 15, no. 3-4, pp. 193–206, 1990.
- [3] W. Chen, Y. Wang, and S. Yang, "Efficient influence maximization in social networks," in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2009, pp. 199–208.
- [4] R. Bellamy, G. Colombo, S. Eshghi, G. de Mel, C. Giammanco, R. Morris, D. G. Rand, L. D. Turner, R. M. Whitaker, and G.-R. Williams, "A computational framework for modelling inter-group behaviour using psychological theory," in *Next-Generation Analyst VI*, vol. 10653. International Society for Optics and Photonics, 2018, p. 106530G.
- [5] B. Latané, "The psychology of social impact," *American psychologist*, vol. 36, no. 4, p. 343, 1981.
- [6] D. Nettle, "Using social impact theory to simulate language change," *Lingua*, vol. 108, no. 2-3, pp. 95–117, 1999.
- [7] J. M. Jackson and B. Latané, "All alone in front of all those people: Stage fright as a function of number and type of co-performers and audience," *Journal of personality and social psychology*, vol. 40, no. 1, p. 73, 1981.
- [8] A. M. Bowen and M. J. Bourgeois, "Attitudes toward lesbian, gay, and bisexual college students: The contribution of pluralistic ignorance, dynamic social impact, and contact theories," *Journal of American College Health*, vol. 50, no. 2, pp. 91–96, 2001.
- [9] N. DiFonzo, M. J. Bourgeois, J. Suls, C. Homan, N. Stupak, B. P. Brooks, D. S. Ross, and P. Bordia, "Rumor clustering, consensus, and polarization: Dynamic social impact and self-organization of hearsay," *Journal of Experimental Social Psychology*, vol. 49, no. 3, pp. 378–399, 2013.
- [10] M. McPherson, L. Smith-Lovin, and J. M. Cook, "Birds of a feather: Homophily in social networks," *Annual review of sociology*, vol. 27, no. 1, pp. 415–444, 2001.
- [11] N. Noë, R. M. Whitaker, M. J. Chorley, and T. V. Pollet, "Birds of a feather locate together? foursquare checkins and personality homophily," *Computers in Human Behavior*, vol. 58, pp. 343–353, 2016.
- [12] N. Noë, R. M. Whitaker, and S. M. Allen, "Personality homophily and geographic distance in facebook," *Cyberpsychology, Behavior, and Social Networking*.
- [13] C. Castellano, M. Marsili, and A. Vespignani, "Nonequilibrium phase transition in a model for social influence," *Physical Review Letters*, vol. 85, no. 16, p. 3536, 2000.
- [14] K. Klemm, V. M. Eguiluz, R. Toral, and M. San Miguel, "Globalization, polarization and cultural drift," *Journal of Economic Dynamics and Control*, vol. 29, no. 1-2, pp. 321–334, 2005.
- [15] C. Gracia-Lázaro, L. Lafuerza, L. Floría, and Y. Moreno, "Residential segregation and cultural dissemination: An Axelrod-Schelling model," *Physical Review E*, vol. 80, no. 4, p. 046123, 2009.
- [16] T. C. Schelling, "Dynamic models of segregation," *Journal of mathematical sociology*, vol. 1, no. 2, pp. 143–186, 1971.
- [17] D. Centola, J. C. Gonzalez-Avella, V. M. Eguiluz, and M. San Miguel, "Homophily, cultural drift, and the co-evolution of cultural groups," *Journal of Conflict Resolution*, vol. 51, no. 6, pp. 905–929, 2007.
- [18] F. Battiston, V. Nicosia, V. Latora, and M. San Miguel, "Layered social influence promotes multiculturalism in the Axelrod model," *Scientific Reports*, vol. 7, no. 1, p. 1809, 2017.
- [19] A. Flache and M. W. Macy, "Local convergence and global diversity: From interpersonal to social influence," *Journal of Conflict Resolution*, vol. 55, no. 6, pp. 970–995, 2011.
- [20] D. Parisi, F. Cecconi, and F. Natale, "Cultural change in spatial environments: the role of cultural assimilation and internal changes in cultures," *Journal of conflict resolution*, vol. 47, no. 2, pp. 163–179, 2003.
- [21] B. Latané, "Dynamic social impact: The creation of culture by communication," *Journal of communication*, vol. 46, no. 4, pp. 13–25, 1996.
- [22] A. Nowak, J. Szamrej, and B. Latané, "From private attitude to public opinion: A dynamic theory of social impact," *Psychological Review*, vol. 97, no. 3, p. 362, 1990.
- [23] J. R. French Jr, "A formal theory of social power," *Psychological review*, vol. 63, no. 3, p. 181, 1956.
- [24] N. A. Christakis and J. H. Fowler, "The spread of obesity in a large social network over 32 years," *New England journal of medicine*, vol. 357, no. 4, pp. 370–379, 2007.
- [25] —, "Social contagion theory: examining dynamic social networks and human behavior," *Statistics in medicine*, vol. 32, no. 4, pp. 556–577, 2013.
- [26] M. S. Granovetter, "The strength of weak ties," in *Social networks*. Elsevier, 1977, pp. 347–367.
- [27] D. Centola and M. Macy, "Complex contagions and the weakness of long ties," *American journal of Sociology*, vol. 113, no. 3, pp. 702–734, 2007.
- [28] R. S. Burt, "Structural holes and good ideas," *American journal of sociology*, vol. 110, no. 2, pp. 349–399, 2004.
- [29] M. T. Hansen, "The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits," *Administrative science quarterly*, vol. 44, no. 1, pp. 82–111, 1999.
- [30] J. F. Fontanari, "Social interaction as a heuristic for combinatorial optimization problems," *Physical Review E*, vol. 82, no. 5, p. 056118, 2010.
- [31] M. Dorigo, E. Bonabeau, and G. Theraulaz, "Ant algorithms and stigmergy," *Future Generation Computer Systems*, vol. 16, no. 8, pp. 851–871, 2000.

- [32] K. Klemm, V. M. Eguíluz, R. Toral, and M. San Miguel, "Nonequilibrium transitions in complex networks: A model of social interaction," *Physical Review E*, vol. 67, no. 2, p. 026120, 2003.
- [33] M. San Miguel, V. M. Eguíluz, R. Toral, and K. Klemm, "Binary and multivariate stochastic models of consensus formation," *Computing in Science & Engineering*, vol. 7, no. 6, pp. 67–73, 2005.
- [34] P. Singh, S. Sreenivasan, B. K. Szymanski, and G. Korniss, "Accelerating consensus on coevolving networks: The effect of committed individuals," *Physical Review E*, vol. 85, no. 4, p. 046104, 2012.
- [35] M. H. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.
- [36] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, "Mixing beliefs among interacting agents," *Advances in Complex Systems*, vol. 3, no. 01n04, pp. 87–98, 2000.
- [37] S. Moscovici, "Toward a theory of conversion behavior," in *Advances in experimental social psychology*. Elsevier, 1980, vol. 13, pp. 209–239.