Semantic networks of simple agent-based models

William Braasch, Pavel Chvykov, Levi Fussell, Xin Ran, Chiara Semenzin*
(Dated: November 1, 2019)

While agent-based models (ABMs) are becoming a dominant tool for studying emergence of social phenomena, they still seem inaccessible far from capturing the immense complexity of social realities. Thus, ever more complicated ABMs are developed, hoping to bridge this gap by better accounting for complexity of individuals. In this work, we illustrate an opposing view, suggesting that complexity can arise from the appropriate choice of collective variables to study, even in the simplest ABMs – shifting the focus from generating the data to interpreting the data. In particular, we simulate a very simple ABM – similar to Axelrod model, but with no homophily – and rather than analyzing the data directly, focus on the semantic network emergent in the culture. We play with how social network topology affects emergence of ideological structures, and entertain questions regarding complex concept formation, "memetic" evolution, and influence of globalization. This way we try to illustrate that a variety of rich sociological questions can be addressed even in a simple model by taking the appropriate perspective.

I. INTRODUCTION

One of the central concepts in complex systems research is emergence: where simple microscopic rules give rise to rich and diverse macroscopic phenomenology [1, 2]. This raises the enticing challenge to try explaining the vast richness of social behavior as an emergent phenomenon. While such ideas have tickled people's imagination for nearly a century [3–5], most of the results in social sciences have largely assumed that the intricacies of society arise from the complexity of individuals. Indeed, it is difficult to imagine that culture, religion, language, science, economics and war could all arise as complex collective phenomena from some simple individual rules. Nonetheless, as social sciences get better at isolating the key factors driving social causation on all scales, we can hope to eventually develop sensible quantitative models of the same.

The main attempts at quantitative understanding of cultural emergence have come from constructing various agent-based models (ABMs) [6]. Perhaps the best-known of these is the Axelrod model [5], which gave rise to a vast number of extensions and modifications, studied both numerically and analytically in the last decades [6]. ABMs as an approach have the mixed blessing of flexibility – it is excessively easy to construct and modify the interaction rules of the agents, which often allows to achieve nearly any collective phenomenon, thus providing little predictive power. This is essentially a generalized problem of over-fitting, which makes the whole approach of using ABMs a bit of an art, with somewhat fuzzy interpretations [7]. Nonetheless, the alternatives are limited, and many of them can be re-stated as certain restricted ABMs, due to the generality of that framework [8].

At the same time, ABMs often seem too simple to faithfully capture social phenomena. The Axelrod model for cultural dynamics, for example, represents each agent as a list of several numbers, which are supposed to stand for "beliefs, attitudes, and behaviors" of individuals, and can then be exchanged among neighbors according to some rules [5]. However, if we hope to really capture the richness of human behaviors, we must admit that representing a culture by just a vector is insufficient. While it is true that an individual may be faithfully represented by some (large) bit-string, it is not purely that information that makes people interesting, but rather how that information shapes their dynamic behavioral patterns in their environment [2]. We must thus look much deeper into the possible perspectives on this information and its consequences for the agents (e.g., [9]).

The goal of this work is to develop an approach to social modeling that addressed these two issues by, on the one hand, keeping the ABMs we use very simple, so as not to drown any insights in a plethora of tuning parameters, while on the other hand, representing the resulting dynamics in ways that have more realistic social interpretation. In particular, we set up a simplified version of the Axelrod model, and look at the resulting data in terms of the semantic networks present in the emergent culture. This way, our algorithm runs on a social network of agents with fixed topology, where dynamical variables ("mental states") live on the nodes. The data we produce, on the other hand, is a semantic network of "ideas", with fixed node identity, but dynamic edge weights and topology, representing how related the two ideas are in the culture. Such semantic networks, while of interest in sociology [10] and computer science [11], have not been explored in ABMs before.

* equal contribution by all authors, order alphabetical

The paper is structured as follows. Section II explains our algorithm and modeling choices. In section III we illustrate the typical dynamics of our model, and show results for some sample questions we could ask in this context. Section IV expounds on the broader takeaways from our setup, showing how it naturally raises new questions not only for numerical models, but for real-world social dynamics themselves.

II. SIMULATION

We designed an agent-based model (ABM) to look at how interactions among individuals give rise to emergence and evolution of ideologies. We construct a social network, whose N vertices represent individual agents, each of whom is described by a length-n binary string (fig.1a and fig.2 left). This string encodes whether or not the given individual knows each of n possible "ideas" that exist in our "society." We further constrain each agent's memory to allow for at-most $m < n$ distinct ideas. The dynamics then proceed asynchronously among random neighbor-pairs, where each time an agent learns one new idea from their neighbor, thus forgetting one old idea of their own due to memory constraint (fig.1). The semantic network is then constructed out of this simulation data (and thereby does not affect the dynamics) by taking the ideas known by anyone in the social network as nodes, and connecting them by edges with weights proportional to the frequency with which the two ideas occur together (fig.2 right). This way it captures the structure and relationships of ideologies present in the culture. Following are the details of the algorithm construction.

A set of interacting individuals can be represented by an undirected unweighted graph, $G = (V, E)$ (left side of fig.2). Individuals are the set of nodes in the graph, where each individual is parametrized by a binary vector of length n, and each individual is further constrained so that the $l1$ magnitude of its representative vector, v_i, is $m < n$: $V = \{v_i \in \{0,1\}^n \ s.t. \ |v| = m \mid 0 < i \leq N\}$. Intuitively, an individual v_i can be thought of as an ordered list of "ideas," where $v_{io} = 1$ represents idea o as present. As there are only so many ideas an individual can have – there are only so many opinions a person can reasonably sustain in their mind – the number of allowable present ideas is limited to $m < n$ for any individual, where n is the total number of ideas in existence. The edges of the graph are the set of allowable interactions of the individual $e : V \times V \to \{0, 1\}$, where individual v_i is capable of interacting with individual v_j iff $e(i, j) = 1$. For the implementation of the simulation, E is an $N \times N$ binary symmetric matrix and V is an $N \times n$ binary matrix.

Ideas are updated by choosing an individual, A, uniformly at random, and then selecting a second 'neighbouring' individual, B, uniformly at random from the set of neighbours $n(A) = \{i \mid e(A, i) = 1, A \neq i\}$ (see fig.1). A 'adopts' a new idea from B by choosing an index o such that $v_{Bo} = 1$ and $v_{Ao} = 0$, and then assigning this idea to v_A, making $v_{Ao} = 1$. Due to the restriction of finite-memory $|v_A| = m$, A must now forget one of their old ideas, and so another index p is chosen such that $v_{Ap} = 1$ and $v_{Bp} = 0$, and is turned to zero. This dynamic represents the diffusion of an idea between two individuals; as two individuals interact more over time, they will adopt ideas from each other becoming more similar. In contrast to Axelrod's cultural model, the probability of individual interaction is independent of their idea-vectors v_i – as such dependence introduces additional complexity and parameters into our ABM, effectively allowing social network topology to change.

The initial topology of the individuals graph has an important effect on the idea dynamics. The connectivity of the individual graph is static over the simulation time, and is defined by the edge matrix. In the next section, we discuss our results for several qualitatively distinct social network topologies.

The semantic network is a reduced graph for understanding the global behaviours of ideas irrespective of the individuals having those ideas (right side of fig.2). It is an undirected weighted graph $H = (D, R)$ where $D = \{d_o \mid 0 < o \leq n, \ \exists v_i \in V \ v_{io} = 1\}$ and edges $r : D \times D \to [0, 1]$. Simply, the semantic network is a graph where a node exists for each idea that is present in *some* individual in G. The weights between the vertices in H are calculated by using the normalised correlation between two ideas in G:

$$r(o, p) = \frac{1}{N} \sum_{i=1}^{N} v_{io} v_{ip} \tag{1}$$

This gives the probability that if you talked to a random individual in the population, they would know both ideas o and p, giving a measure of how closely related those two ideas are.

III. RESULTS

In this section, we look at how the topology of the social network influences the dynamics of the semantic network emergent in the culture. Figure 2 shows snapshots of the two dual networks at some interme-

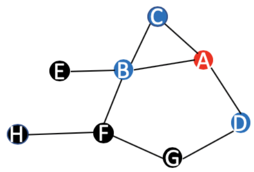

(a) Social network

Timestamp	Agent	Neighbor
t = n	{0, 1, 0, 0, 0, 1, 1}	{1, 0, 1, 0, 1, 0, 0}
t = n + 1	{1, 1, 0, 0, 0, 0, 1}	{1, 0, 1, 0, 1, 0, 0}

(b) Pairwise interaction

FIG. 1: (a). The social network consisting of individuals. Here we focus on node A colored in red and its neighbors colored in blue. (b). Table showing the pairwise interaction between the agent (node A) and one of its randomly chosen neighbor (node B). Binary bits represent their vector of ideas at a certain time. At time t = n, both shaded and dashed squares show the ideas agent A does not possess while its neighbor B does, and the dashed squares show the randomly chosen bit that agent A will learn from its neighbor B through interaction based on the difference in their idea set. And the shaded circle shows the randomly chosen bit which A will lose by obtaining one more idea from B given the limited number of ideas one can possess. At time t = n + 1, the first idea in the vector of A has flipped to 1 as a result of the pairwise interaction.

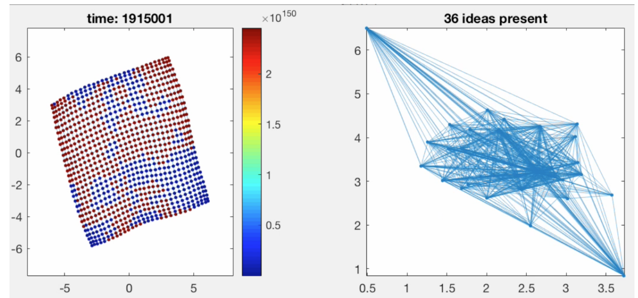

(a) Square lattice social network

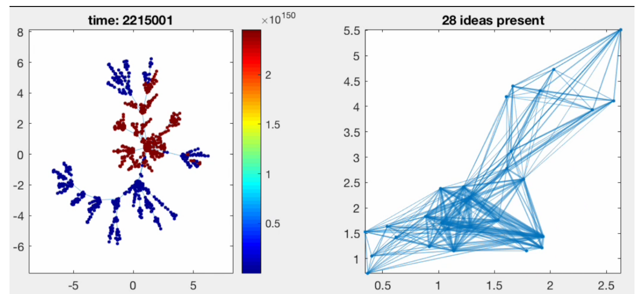

(b) Barabasi-Albert scale-free social network

FIG. 2: Snapshots of the simulation for two different social network topologies at some intermediate times. On the left we show the social network of $N = 900$ agents, with each agent colored according to its individual culture (showing emergence of local communities). On the right we have the emergent semantic network, where nodes are distinct ideas (always more than $m = 7$ – memory capacity, and fewer than $n = 40$ – total existing ideas), and edge weights representing the number of agents who know both of the connected ideas.

diate times in the evolution, for two different social network topologies. The simulation described above always leads to convergence to mono-culture after sufficiently long times, as differences among neighbors get smoothed out by the dynamics. As this happens, and one by one ideas get entirely "forgotten" in the society, the corresponding nodes get disconnected from the semantic network. Figure 3a shows how the size of the semantic network thus drops off over time until reaching its smallest configuration at mono-culture. At this point every agent knows the same $m = 7$ ideas, the semantic network is fully and homogeneously connected, and dynamics stop.

We thus expect all the interesting properties of our semantic network to arise during the transient process at intermediate times: fig. 3a shows that social-network topology has a systematic effect on convergence time-scale. Moreover, the right side of fig.2 shows a qualitative difference between the emergent semantic networks for different social graphs: the scale-free social network gives rise to separated ideological communities, whereas the more intertwined lattice network produces a single interconnected culture. To quantify this distinction, we thought to look at efficiency of the semantic network. However, as the edge weights are crucial here, and efficiency ignores these, we devised a different measure:

Imagine that you walked into this society with idea o in mind. You talk to the first person you see, and ask if they know about o. If they do, you have a conversation, and in the course of which they tell you about idea p. You then go on to talk to another stranger, asking now about p, and if they know it, they will tell you a new idea. We can thus calculate the probability of thus getting from o to o' in any number of steps. Averaging this over all pairs o, o', we get a scalar quantity, which here we call efficiency because it is conceptually related. Numerically, if R is our semantic adjacency matrix, then this number is just $\xi = \langle \sum_{a=1}^{\infty} R^a \rangle = \langle R\,(1-R)^{-1} \rangle$, averaged

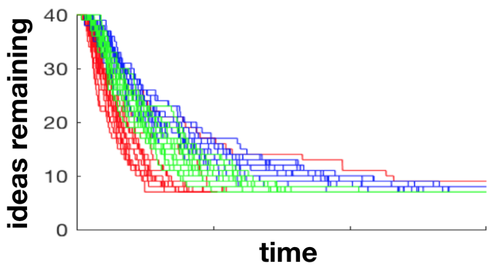

(a) Semantic network size over time

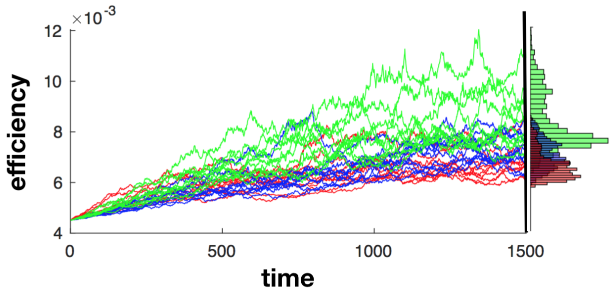

(b) Semantic network efficiency over time, at fixed size

FIG. 3: Simulation results ($N = 900$, $m = 7$, $n = 40$) for several random realizations of three distinct social-network topologies: Barabasi-Albert scale-free network in red (fig.2b), Watts-Strogatz small-world network in green, and square lattice in blue (fig.2a). To get a fair comparison of network efficiencies in panel (b), we restrict the semantic network size to be fixed at 30 ideas (see text). Bar graph on the right shows histograms of the steady-state efficiencies for the three networks.

over all the matrix elements – thus is what we plot in fig.3b. Furthermore, to allow for a fair comparison of semantic network topologies, we must restrict to comparing networks of same size – as otherwise the size will dominate any effect we see. Since semantic network size decays over time, we modified our algorithm to fix a lower bound on the number of ideas – if too many ideas are forgotten, we introduce a random idea into the system, thus keeping semantic network at a constant size, while allowing the social dynamics to run their course. This finally gives fig.3b, showing that the small-world social network produces most tightly connected ideological communities, while the BA scale-free networks produce more disconnected ideologies – in line with the qualitative observation from fig.2.

IV. DISCUSSION

The main contribution of this work is to show that even for very simple dynamical rules, rich structures can arise collectively when we interpret the resulting dynamics from a new perspectives. It seems plausible that at least some of the complexity and structure found in real society is due to the way we interpret our dynamical variables, rather than to the complexity of the interactions themselves. The immediate question this insight raises is how to choose which are the “true” dynamical variables and which are “interpreted” ones? It is exciting to note that this abstract question can be concretely addressed in our model: here it boils down to whether the semantic network dynamics are self-contained. I.e., is the state of the semantic network at time $t = n$ sufficient to determine its state at time $t = n+1$? While it is easy to see that the answer to this question is no, weaker claims could still be exciting: perhaps the history of semantic net over time-period $t \in (n-\eta, n)$ is sufficient to determine the probability over possible configurations at time $t = n+1$. If there is a way to write down the semantic network dynamics in a way without referring to the social network, then we could view the ideas, rather than individuals, as fundamental building blocks of society – agents being merely a carrying medium. Such changes of perspective are integral to many theories in physics, and so may come in handy in studying society as well. This would also put the idea of ”memetics” on a firm quantitative foundation – now truly allowing us to apply tools from ecology to evolution of cultures.

Further development of our model thus has two natural paths: modifying the dynamical algorithm, and finding new perspectives to look at the data. One modification to the algorithm that has been explored is the addition of a decay time for every agent’s ideas. The implementation of decay time has yielded only preliminary results showing a convergence to mono-culture, without any obvious difference to the baseline model. Despite this, further work is required in order to establish the impact of such additional feature, motivated by a more realistic and human-like interaction, where an idea is replaced by another after enough exposure to it, rather than immediately after coming in contact with it.

Another natural modification comes from the observation that membership in modern social networks is an increasingly nebulous topic as virtual interactions increasingly compliment and possibly replace physical social interactions. With the deluge of digital information people now encounter, it is likely that we slowly and subconsciously change our opinions on certain matters that we receive constant exposure to without some watershed event or interaction occurring. This effect could be accounted for via a feedback of information from the semantic network upon the social network. For instance, at each timestep, an action for the chosen agent could be added where an idea not currently held is uniformly selected and then adopted with some probability proportional to the distance on the semantic network between the selected idea and the currently held ideas of the agent. In this way, the agents feel the structure of the semantic network and the forces of persistent, yet subliminal, interactions with common idea and correlated ideas.

Another area of possible investigation is how networks with coevolving agent states and topology, such as the Axelrod model, effect the dual semantic network. This could be applied to current polarizing political environments to explore how echo-chamber dynamics can be disincentivised.

Finally, the other natural direction to take this is comparison to real-world data. Due to the generality and fluidity of our setup, this can be attempted in many different contexts: twitter data, survey results, citation graphs, etc. Something along the lines of the results presented here – a systematic effect of social-net topology on semantic-net efficiency – may be directly tested in such contexts. This may be the next necessary step to develop this work to publication.

[1] Gregoire Nicolis and Cathy Nicolis. *Foundations of complex systems: emergence, information and predicition*. World Scientific, 2012.

[2] Peter A Corning. The re-emergence of emergence: A venerable concept in search of a theory. *Complexity*, 7(6):18–30, 2002.

[3] Nota di Ettore Majorana. Il valore delle leggi statistiche nella fisica e nelle scienze sociali. *Scientia*, 36:58–66, 1942.

[4] Peter Clifford and Aidan Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973.

[5] Robert Axelrod. The dissemination of culture: A model with local convergence and global polarization. *Journal of conflict resolution*, 41(2):203–226, 1997.

[6] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Reviews of modern physics*, 81(2):591, 2009.

[7] Franziska Klügl. A validation methodology for agent-based simulations. In *Proceedings of the 2008 ACM symposium on Applied computing*, pages 39–43. ACM, 2008.

[8] Michael W Macy and Andreas Flache. Social dynamics from the bottom up: Agent-based models of social interaction. *The Oxford handbook of analytical sociology*, pages 245–268, 2009.

[9] Artemy Kolchinsky and David H Wolpert. Semantic information, autonomous agency and non-equilibrium statistical physics. *Interface focus*, 8(6):20180041, 2018.

[10] Jonas Dalege, Denny Borsboom, Frenk van Harreveld, Helma van den Berg, Mark Conner, and Han LJ van der Maas. Toward a formalized account of attitudes: The causal attitude network (can) model. *Psychological review*, 123(1):2, 2016.

[11] Peter Mika. Social networks and the semantic web. In *Proceedings of the 2004 ieee/wic/acm international conference on web intelligence*, pages 285–291. IEEE Computer Society, 2004.