HAL
archives-ouvertes.fr

Emergent Communities in Socio-cognitive Networks

Vincent Danos, Ricardo Honorato-Zimmer

▶ To cite this version:

HAL Id: hal-01976443
https://hal.archives-ouvertes.fr/hal-01976443

Submitted on 10 Jan 2019

Emergent Communities in Socio-Cognitive Networks (revised version)

Vincent Danos
Département d'Informatique
École Normale Supérieure

Ricardo Honorato-Zimmer
School of Informatics
University of Edinburgh

January 26, 2017

Abstract

We investigate a recent network model [13] which combines social and cognitive features. Each node in the social network holds a (possibly different) cognitive network that represent its beliefs. In this internal cognitive network a node denotes a concept and a link indicates whether the two linked concepts are taken to be of a similar or opposite nature. We show how these networks naturally organise into communities and use this to develop a method that detects communities in social networks. How they organise depends on the social structure and the ratio between the cognitive and social forces driving the propagation of beliefs.

1 Introduction

Understanding the mechanisms by which networks self-organise is a major challenge in all information networks: digital, social, and biological. Here we explore a simple (variant of a) model of belief propagation in social networks, originally introduced in Ref. [13]. Our specific question is whether the ways in which belief flows can be revealing of communities in a social network.

The model incorporates a social network formed by individuals (nodes) and social connections through which an exchange of beliefs is possible. In addition, each individual holds its own cognitive network. Cognitive networks consist of a fixed set of concepts (nodes) and relations between them (links), which we call beliefs. Beliefs can be either positive or negative depending on whether the two concepts are taken to be of a similar or opposite nature. Such a two-layered socio-cognitive network of networks is illustrated in Fig. 1.

Whenever there is a cycle in our cognitive network we have an opportunity to check the consistency of our beliefs: if by going around the cycle we end up realising that we believe a concept is of an opposite nature to itself, we have found a contradiction. These contradictions

1

Figure 1: A 2-layer model: on the left a social network is shown. Each node in the social network contains a cognitive network. Here the cognitive network of one node is shown on the right: nodes are concepts and links are beliefs. Beliefs can be either positive (blue) or negative (red).

and how humans deal with them is explained by the theory of cognitive dissonance [7]. In particular, it has been observed that individuals reduce cognitive dissonance and avoid new information that would increase it. In the model, nodes try to minimise such contradictions while maximising consensus.

Our main idea is the following. The socio-cognitive structure of the network favours the emergence of an endogenous and natural notion of trust which can be used to decompose the graph into communities. Individuals distrust others that have a wildly different take on the world. The more different the two world views of two individuals are, the more they will doubt each others' conclusions. Trust in this context is determined by how similar people's beliefs are. We base our community analysis on the idea that communities materialise by how often people find each other agreeing. Your friends are the people you end up sharing beliefs with the most, so to speak.

Concretely, this means that we run repeated randomised belief propagation experiments on top of our static graph and measure how often nodes end up in the same belief group. As we average over initial conditions and look at long term behaviour, per necessity, what we observe is a property of the underpinning network. Results will depend on the "temperature" of the algorithms (how effectively ergodic the underlying Markov chain is), and the relative strength of the social and cognitive forces. The stronger the social field the coarser-grained the community decomposition. Here we will fix temperature (based on empirical convergence times) and study communities as a one-parameter family of average agreement matrices over the population of nodes. This parameterised analysis is reminiscent of persistent homology

techniques in data-analysis [3], and should be resilient to noisy data were we to apply it to real acquaintance networks. We apply our method to the legendary karate club example [14] and find a very convincing decomposition of the club in increasingly fine-grained factions.

1.1 Related work

Many different algorithms for community detection have been created in the past [8]. They are based on markedly different techniques and notions of communities. The majority of them look only at the connectivity of the graph whereas ours takes into account additional structure. Our method is in that sense similar to those based on spin models [12] or coupled oscillators [1]. Here we compare our results in §3.1 to those obtained by the method of Girvan and Newman [11] which produces a similar type of community decomposition. This method uses the *betweenness centrality* of edges, a measure of how many times an edge is used in a minimal path between two nodes, to find a nested decomposition of communities within communities. The results are relatively similar while showing interesting differences.

2 Socio-cognitive systems

A *socio-cognitive network* g consists of finite sets of *individuals* and *concepts*, N and M, a symmetric *link* relation λ on N, and a map $\gamma : N \to \{-1, 0, 1\}^{M \times M}$ assigning a matrix to each individual. This matrix can be seen as a signed graph where edges can be either positive or negative and represents the cognitive network of an individual (as in Fig. 1). Given a node $n \in N$ and concepts $i, j \in M$ we simply write n_{ij} for $\gamma(n)_{ij}$, the belief between i and j in n. The set of all socio-cognitive networks is denoted by $\mathcal{G}$ and those with the same individuals, concepts and link structure by $\mathcal{G}_{N,M,\lambda}$.

For the dynamics, we only allow transitions that change a belief. No transitions can add or remove social connections or individuals. We change one belief at a time and any change of belief is possible. This is a modification from the original model in which only those beliefs that were in disagreement with a neighbour's belief could be changed. This is similar to the voter model [5], and, indeed the model studied here can be thought of as a smoother version of the voter model (where unstructured beliefs battle out for supremacy on the network). This slight change makes the dynamics of our model ergodic on $\mathcal{G}_{N,M,\lambda}$, even though in practice the kinetics might be very "glassy".

We define an energy function to drive the system's evolution. (As usual lower energies represent more favourable states.) This energy function will be composed of a social and a cognitive part.

The *cognitive energy* $E_c : \mathcal{G} \to \mathbb{R}$ describes the cognitive consonance force that is trying to minimise everyone's internal contradictions. Here we only consider contradictions in cycles of length 3, i.e. triangles.

We use balance theory [10] to define the cognitive consonance of a node as

$$C(n) = \sum_{i,j,k \in M} n_{ij} n_{jk} n_{ki}$$

Whenever there is an odd number of negative beliefs in a triangle, the term $n_{ij} n_{jk} n_{ki}$ will be -1. If there is an even number it will be 1. If the three concepts are not forming a triangle because they are not connected, it will be 0.

The total cognitive energy is then

$$E_c(g) = -J \sum_{n \in N} C(n)$$

where J is the strength of the cognitive consonance force acting on the network.

The *social energy* $E_s : \mathcal{G} \to \mathbb{R}$, on the other hand, is determined by trust. We define the trust $T(n, m)$ between two individuals $n, m \in N$ as the similarity of their cognitive networks.

$$T(n, m) = \sum_{i,j \in M} n_{ij} m_{ij}$$

The term $n_{ij} m_{ij}$ will be 1 if the two socially-connected individuals hold the same belief, -1 if they hold the opposite belief, and 0 if any of them do not have a belief between these two concepts. Then the social energy is defined as

$$E_s(g) = -I \sum_{\langle n,m \rangle \in \lambda} T(n, m)$$

with I the strength of the social force. The total energy of the system is then

$$E(g) = E_c(g) + E_s(g)$$

It is easy to see that the minimum energy E_0 that a system can ever reach is obtained when everyone holds the same optimal cognitive network:

$$-E_0 = \binom{|M|}{3} |N| J + \binom{|M|}{2} |\lambda| I \tag{1}$$

Hence a key driver of the dynamics will be the 'soc-cog' ratio I/J, or rather once put in scale invariant form:

$$|M| \cdot \frac{I}{J} \cdot \frac{|\lambda|}{|N|}$$

where $\frac{|\lambda|}{|N|}$ is the average degree (which is preserved by the dynamics).

Clearly, M controls the absolute diversity of opinions in minimum energy states. Specifically, there are $2^{|M|-1}$ such perfectly ordered states. This is one key difference with plain propagation or voter models (on static graphs) where perfect states are far less numerous. The other key difference is less easy to describe and has to do with the ruggedness of the cognitive energy landscape. This is a marked difference to the Axelrod type of belief propagation [2].

4

2.1 Simulation code

Our socio-cognitive systems can be readily simulated using a Metropolis-Hastings algorithm [9]. The code used to run the simulations reported in this paper can be found at `https://github.com/rhz/soccog/tree/gh-pages`. We have also build a web application which runs simulations directly in a browser: `https://rhz.github.io/soccog/`. (It can be advantageous to use a browser which implements javascript efficiently.)

3 Community detection

With our model and simulator in place, we can attack our question. We wish to partition a network into communities by grouping together individuals with the same beliefs (i.e. same cognitive network). The claim in contention is that this extraneous socio-cognitive structure will offer a view of communities parameterised by the soc-cog ratio. Note that for each batch of numerical experiment, we record the *probability* that any two individuals are in the same community at equilibrium (practically defined as acceptance rate steadily under 0.1%). Hence our notion of being in a community is real-valued.

3.1 Karate club

A simple example that has been used many times as a proof of concept for community detection algorithms is the social network studied by Zachary [14] (see Fig. 2 taken from Ref. [6, Chap. 1]). This network maps the friendships within a karate club of 34 people. The fission into two groups that became independent karate clubs was reported in the study and it happened along a faultline that community detection algorithms should be able to identify. This faultline can be intuitively perceived in Fig. 2 due to the way in which the nodes have been arranged. The two groups are

$$\{1, 2, 3, 4, 5, 6, 7, 8, 11, 12, 13, 14, 17, 18, 20, 22\} \text{ and}$$
$$\{9, 10, 15, 16, 19, 21, 23, 24, 25, 26, 27, 28, 29, 30, 31,$$
$$32, 33, 34\}$$

We compute the probability of pairs of individuals ending up with the same cognitive network by running 10,000 simulations and waiting until the rejection rate goes over 99.9% in a window of 100,000 steps. The results are displayed in Fig. 3.

To get a better visual representation of the community structure, we construct a dendrogram showing at which probabilities communities merge. Let $P(n, m)$ be the probability that nodes n and m are in the same community. We start by having each node in its own community and join two communities c, d at probability $p = \max \{P(n, m) \mid n \in c \wedge m \in d\}$. The dendrogram for the karate club network is shown in Fig. 4 and recapitulates nicely the visual intuition of the future split, while identifying other possible splits, e.g. $\{12\}$, and $\{5, 6, 7, 11, 17\}$.

Figure 2: Social network in the karate club studied by Zachary. This figure has been taken from Ref. [6, Chap. 1].

Our results for the karate club network agree to a reasonable extent with the results obtained using the method developed by Girvan and Newman [11]. Interestingly, their method also detects $\{5, 6, 7, 11, 17\}$ as a strong community. However, they see this community join that of node 1 before that community joins the community of node 34. Given the nature of their method, bigger communities tend to disconnect earlier than smaller ones because the betweenness centrality of the edges that connect a peripheral community to the rest of the network are proportional to its size. Our method does not suffer from this bias.

3.2   Random networks

The strengths of the consensus and cognitive consonance forces, I and J, undoubtedly play a role in the formation of communities. To assess it, we look at the communities that arise in random socio-cognitive networks under different conditions. In particular, we first look at the distribution of community sizes. Given g a state of the system, we can define $s_n(g)$ to be the number of communities of size n in g. We have $\sum_n n s_n(g) = |N|$. By running r simulations we obtain a collection G of states at equilibrium. Summing over them we get

$$\sum_{g \in G} \sum_n n s_n(g) = |N|\, r$$

Hence:

$$p(n) = \sum_{g \in G} n s_n(g) / (|N|\, r)$$

is a probability distribution on N. Intuitively, this is (an estimate of) the probability that a random node in N belongs to a community of size n at steady state. We use it as a visual proxy for the state of agglomeration of the nodes in communities.

It is important to realise that changing I and J might have entropic implications. For instance, when doubling the value of both, the shape of the valleys and mountains in the energy landscape will remain constant but the valleys will be deeper and the mountains higher. In other words, I and J together determine the temperature of the system (or equivalently, a change in temperature will produce a simultaneous change in I and J). As a proxy to this intrinsic temperature we look at E_0, the depth of the deepest valley. To *avoid* unintended entropic effects, we keep E_0 constant while we vary $\log_M(I/J)$. The values for I and J are then computed using equation 1.

We construct random networks of $|N| = 100$ individuals with $|\lambda| = 250$ social connections, and assign a random cognitive network of $|M| = 10$ concepts to each individual, and run $r = 10,000$ simulations as described in the previous section. Then we plot the histogram of $p(n)$ versus n for different values of $\log_M(I/J)$ (at constant E_0). A slideshow video of them can be found at `https://tardis.ed.ac.uk/~rhz/hg.mpg`.

One sees that when J is 10 times bigger than I, almost all communities have size 1, that is to say no-one shares the same set of beliefs (and all the individual belief systems are perfectly consonant). This is maximal diversity. (Of course an M which is too small could

prevent this from happening. See the discussion below on the difference between this model and the voter models with few possible opinions.)

As the social field becomes stronger, bigger communities start appearing, unsurprisingly. When I is 10 times bigger than J, it is most likely that a community has a size between 90 and 100. In between, in the transition zone where I, J are of the same order, we see that a large number of different community sizes are possible and equally likely. This raises the question of whether the parameters of the simulations are identifiable even in synthetic data. Indeed the best way to think of our analysis is in terms of a one-parameter *family* of community decompositions. No single value of the soc-cog ratio perfectly summarises the decomposition.

In the previous histogram the correlations between community sizes are lost. So we plot the size of the two biggest communities against their probability of co-occurrence. When it is very cognitive or very social, the sizes of the two biggest communities are concentrated in a small region. Instead during the transition nearly every combination is likely. This confirms that inference of parameters would be difficult: given the community sizes of a series of experiments we might only be able to tell in which of these three regimes we are, but the exact values for I and J would be impossible to infer. The resulting slideshow video can be found at `https://tardis.ed.ac.uk/~rhz/hm.mpg`. A version where the colour range has been truncated to be able to clearly see what happens during the transition can be found at `https://tardis.ed.ac.uk/~rhz/hm-truncated.mpg`. One can see a 'crest line' appearing despite the rather dispersed distributions along which the dynamics is transitioning as we alter the soc-cog ratio. Six of the plots taken during the transition are shown in Fig. 5.

4   Conclusions

There are multiple avenues for further investigation. One is to look for adequate scaling limits that will, in some regimes, admit for simpler approximate formulations of the dynamics and could shed some light on the transition behaviour. Another is to look for efficiency of the simulation. Existing techniques in the domain of energy-based stochastic graph-rewriting [4] should apply. On the applicative side, it would be interesting to look for data on actual social networks and ask how one can work around the difficulty of inferring the key parameters I, J which we have discussed. Which data would be adequate remains to be seen. Perhaps the new Facebook interaction structures with a graded alphabet of likes/unlikes could be used to evaluate belief propagation. One problem is that it is not easy to discriminate between belief acquisition (I change my opinion because of you) and belief confirmation (I am already of the same opinion as you). But that distinction itself might not be of great import.

5　Acknowledgements

References

[1] A. Arenas, A. Díaz-Guilera, and C. J. Pérez-Vicente. Synchronization reveals topological scales in complex networks. *Physical review letters*, 96(11):114102, 2006.

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

[3] G. Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308, 2009.

[4] V. Danos, R. Harmer, and R. Honorato-Zimmer. Thermodynamic graph-rewriting. In *International Conference on Concurrency Theory*, pages 380–394. Springer, 2013.

[5] R. Durrett, J. P. Gleeson, A. L. Lloyd, P. J. Mucha, F. Shi, D. Sivakoff, J. E. Socolar, and C. Varghese. Graph fission in an evolving voter model. *Proceedings of the National Academy of Sciences*, 109(10):3682–3687, 2012.

[6] D. Easley and J. Kleinberg. *Networks, crowds, and markets: Reasoning about a highly connected world.* Cambridge University Press, 2010.

[7] L. Festinger. *A theory of cognitive dissonance.* Stanford university press, 1962.

[8] S. Fortunato. Community detection in graphs. *Physics reports*, 486(3):75–174, 2010.

[9] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.

[10] N. P. Hummon and P. Doreian. Some dynamics of social balance processes: bringing heider back into balance theory. *Social Networks*, 25(1):17–49, 2003.

[11] M. E. Newman and M. Girvan. Finding and evaluating community structure in networks. *Physical review E*, 69(2):026113, 2004.

[12] J. Reichardt and S. Bornholdt. Detecting fuzzy community structures in complex networks with a potts model. *Physical Review Letters*, 93(21):218701, 2004.

[13] N. Rodriguez, J. Bollen, and Y.-Y. Ahn. Collective dynamics of belief evolution under cognitive coherence and social conformity. *arXiv preprint arXiv:1509.01502*, 2015.

[14] W. W. Zachary. An information flow model for conflict and fission in small groups. *Journal of anthropological research*, pages 452–473, 1977.

Figure 3: Heat map showing the probability that two nodes end up in the same community in the karate club network.


Figure 4: Dendrogram showing the probabilities at which communities merge in the karate club network.

10

Figure 5: Heat map showing the size of the two biggest communities at equilibrium for six different values of $\log_M(I/J)$: -0.2 (top left), -0.1 (top centre), 0 (top right), 0.1 (bottom left), 0.3 (bottom centre), and 1 (bottom right). We see a shift from a polarised configuration to one where only one opinion predominates.