Modelling the Dynamics of Collective Cognition: A Network-Based Approach to Socially-Mediated Cognitive Change

Paul R. Smart*††, Winston Sieck†, Dave Braines‡,
Trung Dong Huynh*, Katia Sycara§, and Nigel R. Shadbolt*

*School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK
†Applied Research Associates, Fairborn, Ohio, 45324-6232, USA
‡Emerging Technology Services, IBM United Kingdom Ltd, Hursley Park, Winchester, Hampshire, SO21 2JN, UK
§Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, 15213-3890, USA

††Corresponding author: ps02v@ecs.soton.ac.uk

***Abstract*—A number of studies in the network science literature have attempted to model the effect of network structure on cognitive state fluctuations in social networks. For the most part, these networks use highly simplified models of both cognitive state and social influence. In order to extend these studies and provide the basis for more complex network science simulations, a model of socially-mediated cognitive change is presented. The model attempts to integrate ideas and concepts from a number of disciplines, most notably psychology, evolutionary biology and complexity science. In the model, cognitive states are modelled as networks of binary variables, each of which indicates an agent's belief in a particular fact. The links between variables represent the 'logical' dependencies between beliefs, and these dependencies are based on an agent's knowledge of the domain to which the beliefs apply. Drawing on the psychological notion of cognitive dissonance, it is further suggested that agents are under internal pressure to adopt highly consistent belief configurations, and this identifies one source of cognitive dynamism in the model. Another source of dynamism derives from the structure of the social network. Here, the existence of network ties creates a dependency between the belief systems of connected agents. Cognitive change in such 'coupled belief systems' is modelled using Kauffman's NK(C) model of co-evolutionary development in biological systems. As a final source of cognitive dynamism, the model incorporates the notion of an aggregate belief system (or cultural model), which represents the dominant set of beliefs associated with specific agent sub-groups. By explicitly incorporating the notion of an aggregate belief system into the model, the model supports the analysis of cognitive state fluctuations at the individual (psychological), social and cultural levels. It also provides the basis for future network science simulations that seek to study the complex interactions between these various levels.**

I. Introduction

The emergence of network science as a scientific discipline has provided new insights into how human thought and action may be influenced by social forces. The study of social networks, for example, has yielded a number of important findings regarding the effect of social network structure on the adoption of specific ideas or technological innovations (see [11]), and such findings have important implications for those who seek to influence the spread of specific beliefs, attitudes and behavioural patterns through a target community. Studies of network-mediated cognitive and behavioural change have, however, tended to make a number of rather simplistic assumptions about the nature of both human cognition and social influence. For example, many studies use social networks that are largely undifferentiated at the nodal (agent) level. This situation is clearly unlike the state-of-affairs in the real world, where human agents may differ on a number of dimensions. In addition, many studies overlook the potential role played by pre-existing cognitive states. This is potentially problematic because it seems likely that individuals will be differentially resistant to social influence depending on their current cognitive outlook (for example, states of high internal cognitive consistency may make an individual relatively invulnerable to social influence). In general, the models developed to study cognitive change in social networks suffer from a number of weaknesses. These include (but are not necessarily limited to) the following:

- **Failure to consider the complexity of real-world cognitive systems:** The kinds of cognitive models seen in many network science studies, such as those exploring the dynamics of belief propagation [3, 10], are simple systems comprising single beliefs about the truth or falsity of some worldly fact. These models diverge from the reality of human cognitive state fluctuations in a number of ways. For example, in the human case, we encounter systems comprising multiple beliefs, and such beliefs are often linked together into complex nexuses of logically-consistent belief configurations (see below for more on this).
- **Failure to consider psychological forces and factors:** It is likely that any realistic model of human cognitive change will have to consider at least some aspects of

human psychology. It is well known from the social psychological literature, for example, that the degree of internal consistency between an individual's beliefs attitudes and values can serve as a significant factor in driving cognitive change [2].

- **Failure to consider individual differences:** Although much work in network science tends to overlook differences at the level of individual agents, choosing instead to create homogeneous networks of rather simple nodal elements, it is clear that many real-world networks are highly heterogeneous in terms of their composition. Human social networks, for example, consist of people who may differ along a number of important social and psychological dimensions (e.g. social status, authority, susceptibility to social influence, and so on), and this highlights one way in which many conventional network studies (which countenance the use of largely undifferentiated nodal elements) may fall short as models of collective cognitive flux in real-world social contexts.

This paper describes a model of cognitive dynamics that seeks to address at least some of the weaknesses of extant social network models. The model specifically addresses the weaknesses listed above, and it does so by integrating ideas and concepts from a rich variety of disciplines, most notably psychology, evolutionary biology and complexity science. The model adopts a multi-level approach to socially-mediated cognitive change. Thus, at the level of individual agents, the model adopts a network-based model of cognition that incorporates the psychological notion of cognitive dissonance [2]. This approach incorporates a level of complexity that is conspicuously absent in most network science studies, and it allows us to model one source of cognitive change as the attempt by individual agents to adopt highly consistent configurations of belief states. At the social level, the model incorporates the effect of social relationships on cognitive dynamics. Here, the proposed mechanisms are based on the notion of 'coupled belief systems' – a notion that seeks to highlight the inter-dependence between the cognitive states of socially-connected agents. Finally, we describe a third source of influence in the model – that of culture. Our notion of culture is drawn from studies using the technique of Cultural Network Analysis (CNA) to create network-based models of culturally-specific belief systems [7, 8]. The inclusion of such 'cultural models' into the multi-level model presented here enables us to examine the potential role of cultural forces and factors in shaping the profile of collective thought and action. Together, the three levels provide the basis for network science simulations that seek to study the complex inter-play between a variety of forces and factors on the collective cognitive dynamics of socially-connected agents. In addition to a description of the model, the current paper describes a number of specific ideas for future research (see Section VI). We also discuss how the multi-level network model might be used to further our understanding of the complex cognitive dynamics associated with military coalition operations (see Section VII)[1].

II. Towards a Naturalistic Model of Socially-Mediated Cognitive Change

Developing an interesting model of network-enabled collective cognition arguably requires two essential ingredients: a healthy dose of psychological realism and a rather liberal spicing of real-world simplification. Any network model clearly needs to make some simplifying assumptions about the real-world; otherwise it loses the elements of computational tractability and explanatory concision that make it useful as an aid to both analysis and comprehension. However, a network model should also avail itself of a degree of psychological realism, particularly when aspects of human psychology can be seen to impact the dynamics of information flow and influence in a social network. Many network simulation studies, such as those by Watts [10], make a number of simplifying assumptions when it comes to our understanding of socially-mediated cognitive change, and it is likely that such simplifications undermine the ecological validity of the observed simulation results. Any new model of socially-mediated cognitive change needs to provide richer and more realistic models of cognition and social influence, but it also needs to preserve the features of existing models that make them attractive for network science simulations.

One way in which the models used by previous studies are unrealistic as models of cognitive flux in social contexts concerns the simplicity of the cognitive system possessed by individual social actors. Thus, consider a study undertaken by Glinton et al [3] into the effect of social network structure on belief propagation in a community of synthetic agents. Glinton et al [3] relied on a simplified cognitive model in which agents entertain single beliefs about the veracity of particular facts (e.g. '*F is true*'). This situation is clearly unlike that encountered in more naturalistic, real-world settings. In most real-world contexts, human agents will possess a variety of beliefs, and the truth or falsity of any particular belief will, in part, be determined by the 'logical' inter-dependencies between the beliefs. To illustrate this point, consider a situation where a human decision-maker is attempting to formulate a comprehensive picture of a situation based on a body of incomplete and uncertain information. Suppose that the individual has already received some information that leads her to form certain beliefs. Now, whenever new information about the situation is received, the new information will be evaluated with respect to those pre-existing beliefs. In cases where a conflict is encountered (i.e. new information contradicts a prior belief), the individual will either have to revise her pre-existing beliefs or discount the newly received information. The point here is simply that whenever an agent entertains multiple beliefs, their background (e.g. causal) knowledge often enables

[1]The current paper aims to provide a high-level overview of the multi-level network model. Due to space restrictions we are unable to describe the full details of the model. Interested readers are referred to Smart et al [9] for a more detailed presentation of the model, including a partial mathematical characterization.

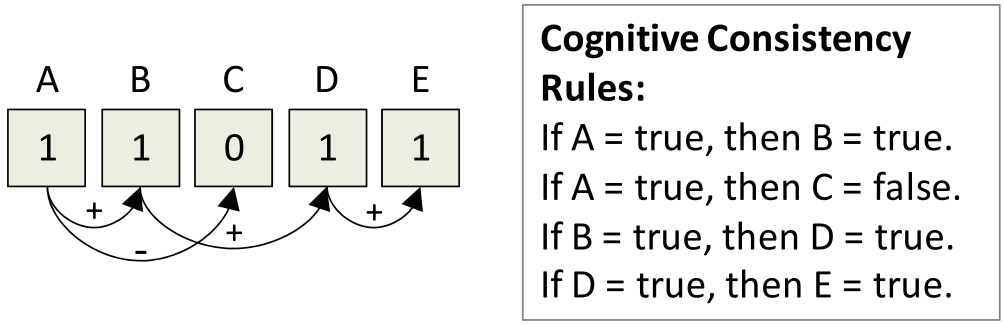

Fig. 1. A primitive belief system exhibiting logical dependencies between specific belief states. For any given belief system, we can enumerate all the configurations (combinations of belief states) that satisfy the consistency rules and are thus cognitively consistent. We can also compute the relative consistency of different configurations based on dependency violations. Note that one extension to this model is to assume that the dependencies are probabilistic rather than all-or-none. For example, we could associate a specific probability with the rule 'If A = true, then B = true' to express the level of certainty or confidence in the dependency expressed by the rule.

them to detect inconsistencies between belief states, and such inconsistencies can influence decisions about whether new information is ultimately accepted or rejected. In cases where conflicting information is accepted, the agent will be forced to revise her existing beliefs (or at least downgrade the certainty assigned to those beliefs); and in cases where such information is rejected, we can begin to understand its lack of influence in terms of an existing complex of inter-dependent belief states.

So one way in which existing models of cognitive flux can be extended is by allowing agents to possess multiple beliefs that are inter-dependent in terms of the extent to which one belief is consistent with another. Such inter-dependencies lead to states of variable cognitive consistency. A state of high cognitive consistency is one in which a set of inter-dependent cognitions (beliefs, values or attitudes, or whatever) are highly compatible and mutually reinforcing. A state of low cognitive consistency, in contrast, is one where we encounter a set of largely incompatible or conflicting cognitions. Suppose that an individual's belief system is represented as a collection of 5 separate beliefs, each of which is modelled as a binary variable (where 1 represents the situation where an agent believes some fact is true, and 0 represents the situation where the agent believes some fact is false). In many cases, we will encounter dependencies between the various beliefs, as determined by the logical or causal structure of the domain to which the beliefs apply. We can represent these dependencies via a set of rules that indicate how a change in one specific belief is likely to influence the other beliefs to which it is connected. Thus, in Fig. 1, we see that if an individual believes that A is true, then it is likely that B is true and C is false. If the individual receives *reliable* information at a later date that C is, in fact, true, then accommodating that information into the individual's belief system would lead to a state of low cognitive consistency. We can see that for any system of beliefs which is modelled in this way we can enumerate all the system states that are consistent (i.e. states in which none of the beliefs violates the dependency rules), and we can also calculate the relative consistency of different system states based on the extent to which the dependency rules are actually violated (see Smart et al [9] for a worked example of cognitive consistency evaluations in the domain of coalition planning).

All this talk of cognitive consistency should strike a chord with those who are familiar with Festinger's [2] theory of cognitive dissonance. Cognitive dissonance theory is a highly popular and influential theory in social psychology, which maintains that individuals are under internal pressure to avoid inconsistent or contradictory cognitions. According to dissonance theory, whenever an individual perceives an inconsistency in their cognitions, they encounter negative mental states (e.g. anxiety, guilt, shame, etc.), and these motivate the individual to attempt to reduce the inconsistency. When people's cognitions are consistent with one another (i.e. high states of cognitive consistency), they are in a state of harmony or cognitive consonance.

There are clearly some important parallels between the foregoing discussion of cognitive consistency and the social psychological notion of cognitive dissonance. Based on the dependencies between a variety of individual belief states, we can calculate the extent to which a particular belief system is in a state of high cognitive consistency (consonance) or low cognitive consistency (dissonance). We can additionally assume that agents will come under internal pressure to avoid dissonant cognitions – they will attempt to avoid states of cognitive dissonance, either by adjusting their specific profile of beliefs (engaging in cognitive change), or by rejecting information that would lead to dissonant cognitions (discounting disconfirming evidence). Given the widespread influence and application of dissonance theory to a variety of areas of psychology, any model of socially-mediated cognitive change is likely to benefit (in terms of psychological realism) by embracing notions of cognitive consistency. This, then, is one way in which the psychological fidelity and plausibility of computational models of socially-mediated cognitive change can be improved. Previous attempts to model the cognitive dynamics of social groups tended to adopt models that made a number of simplifying assumptions about the complexity of human cognition. Such assumptions were clearly necessary in order to make initial inroads into what is, in fact, a highly difficult and complex problem. However, if we are to further our understanding of how the collective cognitive profile of a group of human agents emerges in response to the vagaries of information flow and influence in a variety of networked situations, then it is important that we factor in some of the real-world features that contribute to the stability (or instability) of cognitive states at the individual level. The notion of cognitive consistency, with its obvious linkages to cognitive dissonance, provides, we suggest, the basis for more ecologically-realistic models of socially-mediated cognitive change.

III. Cognitive Consistency, Epistasis and Boolean Networks

We saw, in the previous section, that the inter-dependencies between the various beliefs of an individual agent can be represented as a network of links that exert positive (reinforcing) and negative (contradictory) influences on the various beliefs to which they are connected (see Fig. 1). The resulting system, comprising a network of inter-linked binary variables, is somewhat similar to the notion of Boolean networks, as described by the complexity theorist Stuart Kauffman [4, 5]. According to Kauffman [4] Boolean networks are:

> "systems of binary variables, each with two possible states of activity (on and off) coupled to one another such that the activity of each element is governed by the prior activity of some elements according to a Boolean switching function" (pg. 182).

We suggest that this notion of Boolean networks captures the essential features of the kind of belief system alluded to in the previous section: it enables us to represent the kind of beliefs a particular agent has (e.g. 'on' indicates belief in a particular fact), and it also enables us to represent the linkages or inter-dependencies between those beliefs. The result is what we refer to as an 'epistatic binary belief system', with the notion of epistasis reflecting the fact that the contribution of any one belief to the overall (cognitive) consistency of the total belief ensemble is based, in part, on the linkages between that belief and other beliefs in the same cognitive system (i.e. the cognitive system of a particular agent). The notion of epistasis is also used by Kauffman to represent the dependency between genes in terms of their fitness contributions to the organism of which they are a part[2]. This is important because it establishes a linkage between the notions of cognitive change and cognitive consistency in the case of our model and the notions of evolutionary change and evolutionary fitness in the case of Kauffman's model. Thus, in formalizing the notion of epistasis, Kauffman [4, 5] has developed a mathematical model that has come to be known as the NK model. The NK model is essentially a model for generating evolutionary fitness landscapes with different degrees of 'ruggedness'. Within the model, there are two parameters: N (which refers to the number of genes) and K (which refers to the number of epistatic linkages between the genes). In terms of our discussion about Boolean networks (and belief systems), N refers to the number of binary elements in the network (the number of beliefs in the belief system) and K refers to the number of linkages between the binary elements (the number of inter-dependencies between the different beliefs). Kauffman [4, 5] shows that by using these two parameters, we can generate a variety of 'tunable fitness landscapes', each of which differs with respect to the ruggedness of their topography. When K = 0, there are no epistatic connections between the elements (genes, beliefs, or whatever), and thus each element is free to vary independently of the others (the fitness contribution of each gene is independent of all others in the same genotype). In this situation, Kauffman shows that the fitness landscape is very smooth and contains a single global optima. For $K > 0$, the fitness landscape becomes increasingly rugged, so that when $K = N - 1$ (the maximum value of K), the landscape is maximally rugged.

By applying Kauffman's NK model to this notion of 'epistatic binary belief systems', we can begin to see that each type of belief system can be associated with a particular kind of 'fitness landscape', one whose topography depends on the number of beliefs in the system (N) and the linkages between those beliefs (K). This fitness landscape provides information about the relative fitness of different belief configurations, with the fitness of each belief configuration being determined by the overall cognitive consistency of the belief ensemble. What we have, in essence, is a way of representing the inter-dependencies between beliefs and a way of representing the fitness of different belief configurations in terms of their cognitive consistency. The high points on the fitness landscape indicate states of high cognitive consistency (to which an individual will typically be attracted); the low points indicate states of low cognitive consistency (from which an individual will typically be repelled). The pursuit of more consistent cognitive states (or fitter solutions), in this case, emulates the individual's efforts to arrive at internally consistent cognitive states – the individual is constantly driven to adopt fitter (and therefore more consistent) variants of their current cognitive state (providing that such states are available, or at least reachable).

IV. Coupled Belief Systems

Thus far we have described how the cognitive dynamics of individual agents can be modelled by appealing to a number of well-established ideas in the complexity science literature. However, the main focus of any model of socially-mediated

[2]It should be clear that the fitness contributions of genes encoding for specific traits or features depends, to a large extent, on whatever other genes are simultaneously present in the same genotype. Thus a gene that encodes for dense and heavy bones may be of benefit to a large, flightless bird, but it is unlikely to be useful in the case of a small, flying one. The main point here is that the fitness contributions of genes are rarely independent of other genes: genes interact in complex ways to co-constrain their respective fitness contributions to the organism of which they are a part.

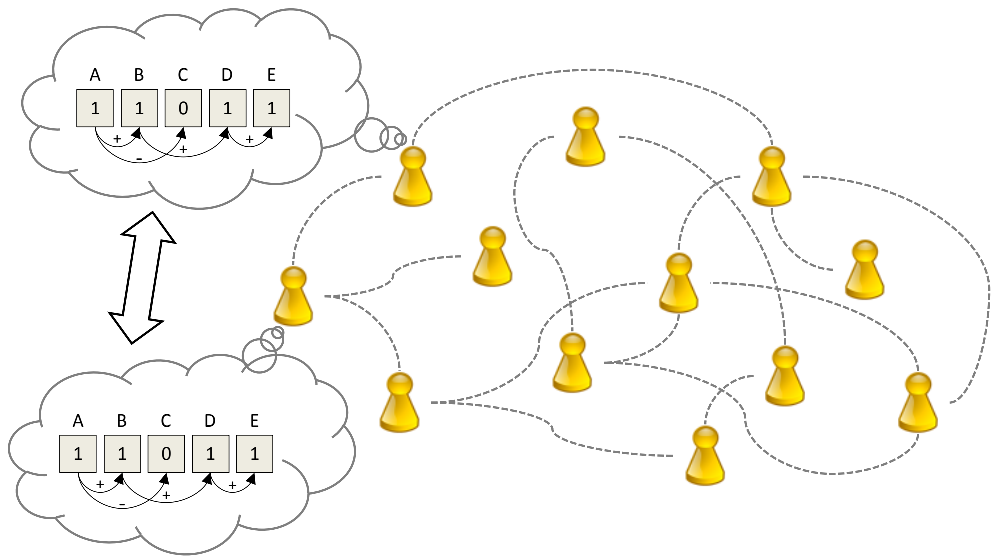

Fig. 2. At the social network level the fitness contributions of individual beliefs are determined by the beliefs of actors to which an agent is connected.

cognitive change is clearly the influence exerted by other elements of the social network (i.e. other agents). Now that we have discussed the intra-individual forces that govern cognitive change, it is time to consider the role of social forces in influencing that change. It is time, in other words, to consider the interplay between network dynamics at both the intra-individual (belief network) and inter-individual (social network) levels.

The approach we have adopted is based on the work of Ahouse et al [1]. Their work is highly relevant to the present discussion because it attempts to extend Kauffman's NK model in order to accommodate the effects of social influence on individual cognition. Ahouse et al's [1] approach is based on an extension to Kauffman's NK model, which is called the NK(C) model. This model was developed by Kauffman to provide insight into co-evolutionary processes. It extends the NK model with an additional parameter, C, which represents the inter-dependencies between the traits of *different* species in terms of their fitness contributions to a given individual. Thus, whereas K represents the linkages or dependencies between the genes of a *particular* individual, C represents the linkages or dependencies between the genes of *different* individuals. This is an important addition to the NK model; it represents the fact that the fitness of any individual is often highly dependent on the features possessed by other individuals. As one individual evolves a particular set of features, it changes the fitness of features possessed by other (inter-linked) individuals. As one species adapts to the exigencies of its environment, so it changes the kind of exigencies that other species must deal with. And a trait whose contribution to evolutionary fitness is irrelevant at one point in time may suddenly become highly significant as the traits of other species change. The image that emerges is thus one of a highly dynamic interdependence between the evolutionary fitness landscapes of different species. The evolutionary landscapes of species are essentially coupled to one another: as one species moves about on its fitness landscape, seeking out locally optimal design solutions, the topographic structure of the landscapes to which that species is connected are systematically deformed. And once the individuals on the newly deformed landscape begin to adapt and explore their new terrain (again searching for local optima), their actions will feed-back to the landscape of the original species, perhaps undermining the integrity of recently scaled peaks and summits.

We can now begin to see how the cognitive influences between agents in a social network might be modelled. Whereas the original application of the NK model (to the notion of intra-individual cognitive consistency and cognitive change) adopted the notion of a fixed fitness landscape against which individuals were motivated to seek out progressively fitter solutions, the image that emerges from a consideration of the NK(C) model is one of a highly dynamic fitness landscape that is in a constant state of flux. The source of such dynamism, in a social network context, is the interdependence between the cognitive states of different social actors. Thus, whether a particular belief configuration is associated with a high fitness score, and thus occupies a lofty peak on a fitness landscape, is not just a function of the consistency of one's own internal cognitions; it also depends on whatever other beliefs are adopted by individuals with whom one interacts. If the individuals within one's social network adopt beliefs that are different from one's own, then one may be influenced to modify those beliefs, particularly if the fitness of one's own beliefs is less than those possessed by one's peers.

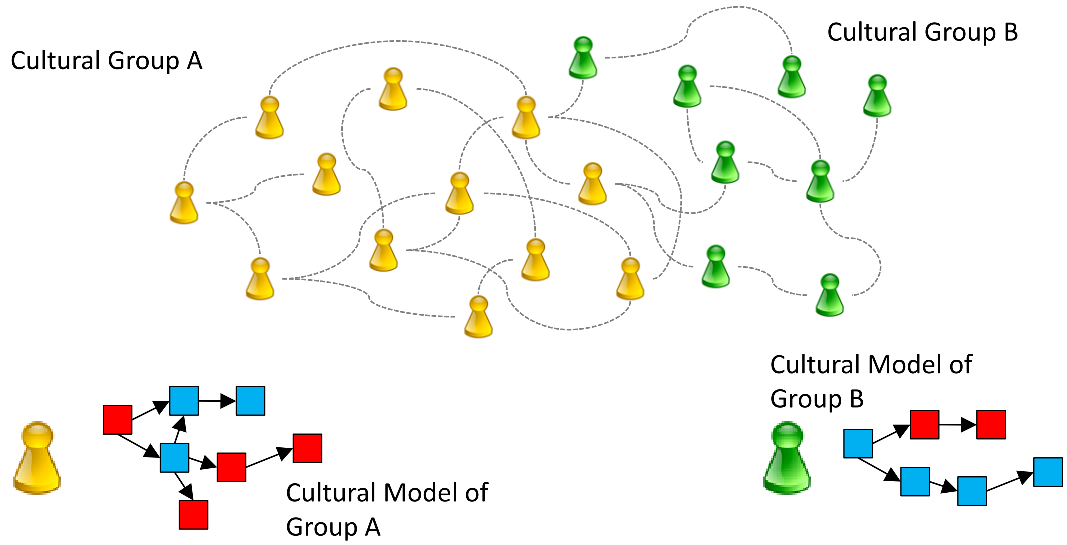

Fig. 3. At the cultural level belief system dynamics are influenced by cultural models. Cultural models act as an additional source of influence on the cognitive dynamics of cultural group members.

The application of the NK(C) model to social networks gives us the notion of what we refer to as 'coupled belief systems'. In place of a single, individual belief system, striving for internal cognitive consistency against a fixed fitness landscape, we now have multiple, dynamic closely-coupled belief systems. As one individual strives for cognitive consistency, so they influence the relative fitness against which other's cognitive states are evaluated (see Fig. 2). And whereas before we were likely to encounter individuals who ended up in highly stable cognitive states (i.e. local optima), and were therefore resistant to further cognitive change, we now have a situation in which the dynamics of intra-individual cognitive change are highly sensitive to the profile of cognitive state fluctuations seen elsewhere in the social network. As one individual changes their beliefs, so they influence the stability of cognitive states in those individuals to whom they are connected via social ties. In essence, the collective beliefs of all individuals in a social network provides the cognitive niche in which the cognitive state of any given individual evolves.

V. Cultural Models and Collective Cognition

The notion of epistatic binary belief systems is intended to provide a computationally-tractable model of belief networks at the level of individual human agents. Such systems need not, however, be restricted to the level of specific individuals. An alternative approach is to develop an aggregate model of beliefs (and their associated linkages), which reflects the coarse-grained regularities associated with the belief systems of *multiple* individuals from a specific population. The technique of Cultural Network Analysis (CNA) is designed to produce such models [7, 8]. CNA is a methodology that synthesizes techniques from a range of disciplines, including cognitive anthropology, cultural psychology, naturalistic decision making, and decision analysis. Its aim is to develop models (referred to as cultural models) that capture the key (cognitive) commonalities among a group of individuals in terms of their shared concepts, beliefs and values. Cultural models reveal the relative differences and similarities between social groups, and they can be used as the basis for discriminating between specific cultural and sub-cultural groupings.

Using the notion of cultural models, we can model the influence exerted by cultural-level forces on cognitive change. The approach we have taken (see [9]) is to cast cultural models as a form of epistatically-linked belief system, similar to that seen in the case of individual agents. However, whereas an epistatic binary belief system at the level of individual agents represents the beliefs (and associated inter-dependencies) of a *particular* agent, the Boolean network projection of a cultural model represents the common elements of the belief systems associated with *multiple* agents, specifically those belonging to the same cultural group (see Fig. 3). The primary significance of this shift away from individual belief systems to more collective, aggregate belief systems is that it enables us to think about the potential for cognitive change mediated by cultural-level influences. For example, we could assess the potential role played by cultural group membership in initiating or preventing cognitive change by representing the cultural model as the belief system of another (abstract) agent in the model (a kind of 'cultural model agent'). We could then establish a strong social tie between the cultural model and agents who are members of the cultural group represented by the cultural model.

The use of cultural models also enables us to study the interactions between the belief systems of different cultural groups. Thus, imagine a situation where multiple cultural groups are manifest in the same social network (see Fig. 3). We could investigate the impact of cross-cultural exposure in this network by (again) representing the cultural models as the belief systems of specific 'cultural agents' and establishing social influence relationships between these belief systems and the belief systems of agents from other cultural groups. Something (roughly) similar to this is proposed by Ahouse et al [1]. They recognize that rather than considering the belief systems of specific human individuals, we could consider the belief systems associated with nation states. Thus, if we consider the general beliefs held by, say, people from the US, Russia and Western Europe, we can develop an NK(C) model in which N represents the beliefs held by the various populations, K represents the dependencies between the beliefs held by each population, and C represents the dependencies between the beliefs held by different nation states (e.g. between US and Russia). As suggested by Ahouse et al [1], these beliefs could be about the relative merits of military spending for security, for the overall economy and for the extent of government involvement in domestic social policies. Inasmuch as these respective belief systems form a coupled belief system, we may expect the cognitive dynamics of one population to influence the fitness landscapes of other, connected, populations. A shift in public opinion in one cultural group may result in, perhaps unexpected, changes in the beliefs and ideologies of other cultural groups.

VI. Research Issues

The model of socially-mediated cognitive change presented here provides the basis for network simulation studies examining the effect of multiple forces and factors on the collective cognitive dynamics of multi-agent systems. A number of interesting research issues arise in the context of such studies. The following are some examples of the kinds of research questions that could be pursued in future empirical studies.

1) **How does the profile of inter-agent interaction at the social network level contribute to the emergence, modification and dissolution of cultural groups?** The aim here is to understand how the dynamics of inter-agent interaction (as perhaps determined by time-variant patterns of functional connectivity at the social network level) contributes to the evolution of distinct cultural models. We can use the methods developed as part of CNA to construct cultural models on the basis of information about the beliefs held by agents in a particular population. In the case of our simulation model, we have complete access to the beliefs of all agents in the simulation, and we can thus automatically construct and modify cultural models based on the changing profile of beliefs possessed by all the agents. We now encounter an interesting set of questions relating to the emergence of distinct cultural models in a multi-agent system. If, for example, we begin the simulation with multiple cultural models (reflecting the presence of multiple cultural groups), how will the cultural models change over time in light of the dynamics of information flow and influence at the social network level? Will all the agents adopt one or other model (reflecting a process of cultural convergence or cultural assimilation), or will the models merge to form a single hybrid cultural model (perhaps reflecting a process of cultural hybridization or 'cultural syncretization')? Perhaps the original cultural models will be preserved, but their boundaries at the social level will be highly porous (i.e. different agents at different times will switch between different cultural communities). Clearly, the combination of social networks, coupled belief systems and automated cultural model construction provides a rich source of simulation opportunities, with potentially profound implications for our understanding of the emergence, modification and dissolution of cultural groups.
2) **How do cultural-level influences affect individual cognitive change?** The aim here is to explore the interaction between forces for cognitive change at the individual, social and cultural levels. One possible simulation involves the creation of additional sources of influence between an individual's beliefs and the beliefs of the cultural group to which they belong (the influence could be represented via the inclusion of additional epistatic linkages, perhaps with higher weightings than those associated with other linkages). The cultural models could still be subject to automatic creation and modification (as discussed above); however, in this case, the cultural models would provide an additional source of influence on the cognitive dynamics of individual agents – agent cognition would be influenced by the aggregate beliefs of the cultural group to which those agents belonged.
3) **Do forces at the inter-cultural level contribute to the cognitive stability of culturally-heterogeneous groups?** The aim here is to examine the notion that forces operating at the cultural level (the level of cultural models) may contribute to a form of 'ideological equilibrium' in which the cognitive outlooks of different cultural groups are stabilized by the emergence of Nash equilibria at the level of cultural models. A detailed discussion of this issue is, unfortunately, beyond the scope of the current paper; however, see Smart et al [9] for an extended discussion of this issue.

VII. Relevance to Military Coalitions

As revealed by previous research (e.g. [6]), military coalitions often consist of culturally-heterogeneous collections of individuals. Such heterogeneity is, of course, not necessarily surprising: military coalitions often bring together individuals with different skills, experience, attitudes and values. Understanding the effect of such cultural differences in terms of the potential for inter-agent collaboration, communication and coordination is a major focus of research interest. Another issue, however, concerns the effect that inter-agent interaction may have on cultural differences themselves. Is it possible, for example, that certain kinds of inter-agent interaction, as mediated by the structure of the coalition communication network, might result in cognitive change in one or more cultural groups? And, if change does occur, what is the nature of this change? Does it reflect a process of cultural convergence, assimilation or hybridization? We have, as yet, little understanding of how social network variables interact with psycho-dynamic and cultural-level influences in coalition settings. The model described here provides one means of identifying some potentially important relationships that may have implications for our understanding of how cognition evolves in social contexts. It also provides the basis for empirical studies that may shed light on how best to develop interventions that promote specific forms of socially-mediated cognitive change.

VIII. Conclusion

A number of studies in the network science literature have attempted to model the effect of network structure on cognitive state fluctuations in social networks. For the most part, these studies rely on models that make a number of rather unrealistic assumptions about the nature of real-world cognition and the mechanisms of social influence. In order to address the rather ersatz nature of these models, we have described a multi-level network model of socially-mediated cognitive change. Cognitive change in this model derives from three sources (roughly equivalent to forces acting at the psychological, social and cultural levels). Firstly, we saw that individuals might come under internal pressure to change their profile of beliefs based on the perception of specific inconsistencies and conflicts. In this case, cognitive change is motivated by a need for consistency between beliefs, and the individual strives to minimize inconsistent cognitions. Secondly, we saw that cognitive change could come about as a result of forces that operate at the level of social networks. In this case, the individual is influenced by the beliefs possessed by other individuals, and cognitive change is likely to strike a balance between social influence and cognitive dissonance. Finally, we saw that there might exist collective forces that influence the cognitive dynamics of individuals within specific cultural groupings. This level of analysis emphasizes the potential role that inter-cultural relationships play in shifting or stabilizing the cognitive outlook of individual group members.

The dynamics of information and flow and influence within and between these various levels is governed by Kauffman's [4, 5] NK(C) model of evolutionary and co-evolutionary dynamics. The NK(C) model provides the algorithmic basis for such simulations, and the equations developed in support of the NK(C) model can be easily adapted for the current case. Such features support one of the key requirements of a network-based model of socially-mediated cognitive change, namely its amenability to formal characterization and computer simulation. A second requirement relates to the ecological plausibility and validity of the model – its ability to adequately capture those features of the real-world that make the model of genuine explanatory interest and predictive relevance. The multi-level network model presented here is a first step towards meeting this demand for ecological validity. It is a significant extension to the models seen in previous work, and it provides a solid basis on which future network science simulations can be undertaken.

Acknowledgment

This research was sponsored by the US Army Research laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. 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 US Army Research Laboratory, the U.S. Government, the UK Ministry of Defence, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

[1] J. J. Ahouse, E. Bruderer, A. Gelover-Santiago, N. Konno, D. Lazer, and S. Veretnik. Reality kisses the neck of speculation: A report from the nkc workgroup. In L. Nadel and D. Stein, editors, *1991 Lectures in Complex Systems, SFI Studies in the Science of Complexity*, volume 4, pages 331–353. Addison-Wesley, Reading, Massachusetts, USA, 1992.

[2] L. Festinger. *A Theory of Cognitive Dissonance*. Stanford University Press, Stanford, USA, 1957.

[3] R. Glinton, P. Paruchuri, P. Scerri, and K. Sycara. Self-organized criticality of belief propagation in large heterogeneous teams. In M. J. Hirsch, P. M. Pardalos, and R. Murphey, editors, *Dynamics of Information Systems: Theory and Applications*, Springer Optimization and Its Applications. Springer, Berlin, Germany, 2010.

[4] S. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York, USA, 1993.

[5] S. Kauffman. *At Home in the Universe: The Search for Laws of Self-Organization and Complexity*. Penguin Books Ltd, London, England, 1995.

[6] L. J. Rasmussen, W. R. Sieck, and P. R. Smart. What is a good plan? cultural variations in expert planners concepts of plan quality. *Journal of Cognitive Engineering and Decision Making*, 3(3):228–253, 2009.

[7] W. R. Sieck and L. Rasmussen. Cultural network analysis. In *1st Annual Conference of the International Technology Alliance (ACITA)*, Maryland, USA, 2007.

[8] W. R. Sieck, L. Rasmussen, and P. R. Smart. Cultural network analysis: A cognitive approach to cultural modeling. In D. Verma, editor, *Network Science for Military Coalition Operations: Information Extraction and Interaction*. IGI Global, Hershey, Pennsylvania, USA., in press.

[9] P. R. Smart, T. D. Huynh, D. Braines, K. Sycara, and N. R. Shadbolt. Collective cognition: Exploring the dynamics of belief propagation and collective problem solving in multi-agent systems. In *1st ITA Workshop on Network-Enabled Cognition: The Contribution of Social and Technological Networks to Human Cognition*, Maryland, USA, 2010.

[10] D. J. Watts. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences of the United States of America*, 99(9):5766, 2002.

[11] D. J. Watts. *Six Degrees: The Science of a Connected Age*. William Heinemann, London, UK, 2003.