A Connectionist Modeling Approach to Rapid Analysis of Emergent Social Cognition Properties in Large-Populations

Kalyan S. Perumalla and Jack C. Schryver

Abstract—**Traditional modeling methodologies, such as those based on rule-based agent modeling, are exhibiting limitations in application to rich behavioral scenarios, especially when applied to large population aggregates. Here, we propose a new modeling methodology based on a well-known "connectionist approach," and articulate its pertinence in new applications of interest. This methodology is designed to address challenges such as speed of model development, model customization, model reuse across disparate geographic/cultural regions, and rapid and incremental updates to models over time.**

I. INTRODUCTION

BEHAVIORAL modeling techniques in vogue, such as rule-based agent models, while being adequate for small scenarios or simple compositional behaviors, are becoming less suitable for large population scenarios with rich and complex individual behaviors. In regard to modeling power, they are fundamentally limited by their inadequacy in capturing latent states and latent relationships among various behavioral states. Additionally, due to reliance on distilled and codified data, such as poll and survey results, traditional approaches are constrained by the quality and speed of collection of such data. Consequently, the resultant modeling methodologies suffer from premature, incomplete or inappropriate abstractions, missing important low-level relationships among latent states that are not explicitly modeled. Here, we identify and propose the use of an alternative modeling methodology based on a "Connectionist Model" framework that has lately begun to populate the social psychology literature, and that holds promise to address several shortcomings of existing approaches in large-population behavioral studies.

A. Motivation

Recently, an increased amount of interest from the academic, government and defense communities is evident in simulation-based behavioral studies for large-population scenarios. In the absence of alternative or better methodologies, there appears to be intense interest in using agent/rule-based approaches to understanding behaviors of real social systems. However, methodological problems are arising when existing rule-based approaches are being applied to the new needs of scale and speed in large-scale scenarios of interest.

To illustrate, the defense and intelligence communities are interested in simulating and understanding social behavioral implications of military alternatives in operations such as pre- or post-conflict campaigns. For understanding the implications, the populations of interest could be either domestic or foreign or both. The scales of interest can range from towns and cities to states and countries as well, although the quantitative nature of analysis is usually progressively relaxed to become more qualitative in nature for larger population bodies. Decisions such as the determination of the amount of foreign aid, or the emphasis on infrastructure (living condition) development in the theater of interest may be based on the expected or afforded level of behavioral benefit (e.g., cooperation or non-hostile disposition) from the subject population. Behavioral modeling and simulation may need to be employed in relatively short order when new theaters of interest are identified, and when planning (e.g., deployments or resource allocations) must be performed relatively quickly. The behavioral analysis problem is exacerbated by the fact that the simulation models may need to be applied to population classes that were not specifically studied in advance. For cost-effectiveness, the same models may be reused for disparate theaters, making it necessary for the models to be sufficiently flexible to be customized and configured for different population characteristics.

Unfortunately, little existing research is available to deal with the *methodological* and *modeling* challenges in rapid analysis of behaviors in large and diverse sets of populations. It has been unclear as to how the models must be structured from the outset to be able to capture the potential diversity of behaviors within a population body, as well as to make the models reusable in some fashion across different disparate population bodies. An understanding of the levels of modeling and simulation efforts imposed by the new needs is missing from a programmatic and methodological viewpoint. New modeling frameworks and methodologies would be useful to develop in order to address the methodological problems by design from the very outset. Here, we identify some of these shortcomings of existing approaches, and propose a direction that

Manuscript received May 20, 2009. This paper has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the U.S. Department of Energy. Accordingly, the United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

The authors are with the Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA. Phone: (865) 241-1315/574-4710. Fax: 865-576-0003. Email: perumallaks@ornl.gov, and schryverjc@ornl.gov.

promises to address the new methodological- and modeling issues and challenges.

B. Background and Related Work

Among existing approaches, one of the most popular is built on agent-based modeling and simulation methodologies. Agent-based simulation models have been used successfully over the past decade or more in efforts to increase understanding of the behavior of social systems. During initial phase of that research, models with simple cognitive units representing individuals were invoked to establish the potential for agent-based models to capture complex systems behavior widely believed to characterize social systems (e.g., [1]). However, it is unclear how to fit existing agent-based approaches into a broader modeling methodology for realistic models of large and diverse sets of populations. For example, when an agent-based model is attempted to be reused from one population to another, the relative importance values of behavioral elements (such as attitudes and preferences) may change dramatically (say, religious affiliation may have a much larger role to play than economic conditions). To incorporate such qualitative changes, the agent-based models typically need to undergo significant coding and structural changes, incurring new modeling and simulation efforts from subject matter experts.

Besides methodological issues, agent-based models appear to also be largely incapable of accurate prediction or forecasting, and are harder to generalize. One important reason for this inadequacy is lack of psychological plausibility. Existing agent-based simulation architectures have been criticized because of reliance on assumptions about the human brain of serial processing and simplistic cognition that ignores parallel, distributed and automated processing. At a minimum, psychologically plausible models should account for social phenomena reported in the social science literature. For example, social psychologists have identified group biases (illusory correlation, accentuation, sub-typing and out-group homogeneity) that are responsible for group stereotyping, which are not easily explained by sequential processing, but can be modeled as parallel cognitive processes.

Here, we identify a "connectionist modeling framework" that lately has been gaining research attention in traditional (small- to medium-sized group) psychological modeling [2-5], and articulate its excellent potential for application to large-sized populations. We present a methodology based on connectionist modeling to address methodological as well as modeling challenges raised by new behavioral modeling and simulation applications in large population scenarios.

C. Organization

The rest of the document is organized as follows. The modeling issues and challenges in new large-scale scenarios are outlined in Section II. A brief introduction to the connectionist modeling frameworks and methodologies is documented in Section III, illustrated with an example. In Section IV, the potential for configuring, customizing, reusing and refining connectionist-based models with non-traditional data sources is documented. The paper is summarized and concluded in Section V.

II. Modeling Methodology: Issues and Challenges

With the new level of scale and detail being explored by social science and social computing communities, it is useful to understand the nature of ensuing challenges. Here, we summarize some of the most important issues and challenges, in terms of modeling *methodology* as well as the framework for *modeling behavior* at an individual-level.

A. Methodological Challenges

A modeling methodology is the process by which models are developed, verified, validated, customized, configured, simulated, interpreted, updated, and reused. The methodological challenges that arise in behavioral modeling for large populations include the following:

- The size of population can be as large as 10^5-10^7 individuals. Such large individual-counts make it necessary to reexamine the methods used for smaller scenarios such as small control groups.
- The models may need to be applied and/or reused across potentially diverse set of populations; for example, across two entirely different countries of interest.
- Ideally, the entire modeling and simulation process will need to be rapidly responsive to calls for new scenarios to be analyzed, and operate so in the potential absence of detailed background behavioral data.
- In case where historical data is available, the modeling methodology must afford automation to assimilate the data to morph into a model.
- In a theater of interest which is evolving relatively rapidly, the modeling framework must be able to incorporate newly arriving, updated information into the underlying model.

B. Behavioral Modeling Challenges

In addition to the methodological challenges, the specific modeling framework must be capable of dealing with the diversity and richness of phenomena that appear in the behavioral classes of interest. This is especially important in order to support the complexity even at the level of each individual's cognitive processes. Some of the important challenges in developing rich behavioral models are:

- The models need to incorporate richness of attitudes and behaviors at the individual level, to account for the potential for immense diversity and latency of states in large population counts.
- The models must account for latent states or potential hidden linkages among attitudes of individuals to objects. Ideally, this capture of latent states and latent linkages must not impose the burden of explicit coding

or manifestation by the modeler, since such explicit coding would involve undue effort for large models and/or multitude of scenarios.

- The models must have less reliance on slow data collection processes such as surveys and polls. Developing and configuring the models with more readily and rapidly available data sources (such as images or financial flow data) can make the models more resilient to real-time demands.
- The models must reduce reliance on data that potentially suffers sampling bias, such as web-based documents, since such data can be statistically infeasible to use in theatres such as third-world countries.

As will be described in greater detail, a connectionist modeling approach is useful in addressing many of these issues and challenges.

III. Connectionist Models

A connectionist approach to the construction of multi-agent-based models offers much promise in moving toward more psychologically plausible models. Several researchers have called for a more complete account of individual cognitive factors in social simulation models [6-8]. Greater cognitive complexity at the individual-level may be achievable, albeit at the expense of slightly increased computational power.

A. *Salient Features*

In a connectionist modeling framework, each individual is modeled as a connectionist network composed of large numbers of simple cognitive units. Connectionism employs models of parallel and distributed computation performed by highly interconnected simple cognitive units. Knowledge is represented in patterns of activation over these units. The units are highly simplified versions of real neuronal circuitry and processing. The framework has been previously shown to be sufficiently powerful to demonstrate how psychological and social phenomena can emerge from connectionist networks of individual cognition. Additional power can be gained by using connectionist networks at both the individual- and interpersonal-levels of analysis, thereby creating the necessary conditions for multiscale compositional behaviors.

Many connectionist models are possible, but all approaches tend to share several basic micro-cognitive elements. According to one of the seminal sources [9], eight major aspects can be identified:

- A set of processing units
- A state of activation
- An output function for each unit
- A pattern of connectivity among units
- A propagation rule to spread activation throughout the network
- An activation rule for combining inputs to a unit
- A learning rule to modify weights on the basis of experience
- An environment.

Among the aforementioned items, the core elements are formed by a graph. The nodes of the graph are the processing units, which are either labeled by a cognitive aspect or unlabeled to account for undetermined intermediate cognitive functions. The edges of the graph are labeled with real-valued weights, with the allowance that the weights can be evolved over time. Units undergo "activations," which prompt them to emit real-valued signals on their output arcs. In general, activations can be spontaneous, timed, or driven by reception of signals on incoming arcs. Input or output arcs' weights can be modified by units over time. Weights are used by receiving units to evaluate an aggregation function over all incoming signals. Note that weights are allowed to be negative, to model cognitive inter-relationships such as dissuasion or counter-influence. Often, a subset of units is identified as "input activators" and another subset as "output units," but such distinction is more out of convenience than out of conceptual necessity.

In contrast to agent-based models that are based on state variables and explicitly-coded rule sets, the connectionist modeling approach blurs the distinction between memory and computation, thus providing a more plausible representation of human cognition from a neuroscience perspective. Computations are performed by highly interconnected simple cognitive units, somewhat analogous to the actual circuitry of the human brain. Units are activated through either internal or external excitation (or inhibition) and spread activation to other units inside the individual's connectionist network according to the strengths of the arc weights. The pattern of activation that develops in response to an external stimulus represents the short-term memory (STM) of the individual network. Long-term memory storage resides in the matrix of arc-weight values. Weight values are temporally stable in the short term, but adapt through application of specific learning rules.

The "delta learning rule" has been very popular among the connectionist models designed by social psychologists; other investigators prefer Hebbian learning. The former learning method allows the internal state of activation to gradually reproduce patterns found in the environment or a teacher, whereas the latter drives units that activate simultaneously to increase their tendency to do so. Many social psychological models are constructed with a recurrent, auto-associator architecture. Although difficult to train, recurrent architectures allow models to exhibit dynamic behavior while the model settles into a stable pattern of activation in response to an external stimulus.

The connectionist modeling approach is amenable to many cognitive processes, but a most natural fit obtains in the representation of partially automated processes such as attitudes, beliefs, stereotypes, biases, and opinions. Mental processes operating outside conscious awareness can be modeled without application of deliberate rule-based coding.

B. *Application of Connectionist Models*

Connectionist models have begun to populate the social psychology literature for over a decade. In particular, various aspects of social cognition have been modeled, including person impression formation, causal attribution, stereotype learning and change, decision making, group identification, group bias and attitude formation and change [10, 11]. For example, one model [12] explains the emergence of group biases from connectionist networks of individual cognition. Group biases emerge because confirmatory evidence gradually strengthens the connection between attribute and object during learning; features that are more diagnostic than others compete more successfully for connection strength in learning. Individual-level models provide a connectionist account of beliefs and attitudes by representing objects/concepts and attributes in interconnected processing units. A 'cognitive trust model' links individual cognitive models together into a social interaction model that explains communication and social influence effects within a connectionist framework.

Connectionist models are not used in social psychology to model neurons in the brain. They approximately model assemblies of neurons at a micro-cognitive level of description that is intermediate between neurons and the symbolic or rule-based level. Phenomena that are typically described at the symbolic level emerge from a connectionist layer that implements general-purpose micro-cognitive processes. Although connectionist networks are mathematically similar to neural networks described in the machine learning literature, they have divergent purposes and semantics. Typically, connectionist networks in the cognitive sciences are not pattern associators or function-approximators. There is typically more clarity regarding the meaning of the components which makes (localist) connectionist networks easier to configure and interpret as compared to standard neural networks.

Connectionism *vs*. symbolic processing has been debated within cognitive science for twenty years. Each approach has many proponents, but connectionism, which has been overlooked in the artificial intelligence community, has not been explored much in multi-agent-based architectures. We believe that connectionist models offer many comparative advantages over symbolic computation. They provide a parallel and mutually constraining approach to human brain computation as an alternative to serial processing architectures. Connectionist models do not require a central executive for as a clearing house for information. Consequently, implicit and automatic mechanisms can be modeled without reference to explicit conscious reasoning. Cognitive phenomena such as biases and heuristics are accounted for not by the introduction of ad hoc cognitive limitations, but as emergent properties of general micro-processes that are otherwise adaptive.

A particular strength of connectionist networks is that they can represent the interdependent nature of the attitudes, beliefs, stereotypes, biases, and opinions. Lack of interdependency among mental processes and representations is a known deficiency of agent-based architectures. For example, persuasive political communications affect beliefs and attitudes toward other issues, even though the message may not directly reference those other issues.

Connectionist models are usually divided into localist [2] and distributed [13] categories. Localist models are distinguished by a fixed semantics over the entire set of nodes for which it is always possible to retrieve the meaning of individual nodes. Distributed modes have a subset of unlabelled nodes. It is not in general possible to determine the meaning of a single node by examining its level of activation. Semantic representation occurs instead through examination of the pattern of activation over the entire set of unlabelled nodes. Both types of models are widely in use within the connectionist community, although localist models are enjoying somewhat of a resurgence in recent years. Localist models can bring greater explanatory power to a model if relevant features are known a priori. This explanatory power comes at the expense of the inability to learn new concepts, which distributed networks are capable of doing.

C. *Intra-individual Connectionist Network*

Van Overwalle and Siebler [3] introduced a connectionist model with a localist representation of attitudes that is useful as an example because of psychological plausibility and transparency. Other similar models of attitude formation and change have been developed but are omitted here for the sake of brevity [14]. The model in [3] was inspired by two well-known theories of attitude formation and change. The model depicts two routes to attitude change according to the Elaboration Likelihood Model of persuasion [15]: central and peripheral. The central route to persuasion can be followed by elaborated discussion of positive and negative attributes of the attitude object. The attitude is constructed in short term memory (STM) using a weighted summation of all activated attributes according to expectancy-value theory introduced by Fishbein and Ajzen [16].

The approach taken by [3] was to apply a recurrent auto-associative architecture with the delta learning rule. The authors were able to demonstrate the generality of the attitude model by successfully reproducing simulated data that mimicked experimental data showing the use of heuristics of length, consensus, expertise and mood.

D. *Inter-individual Connectionist Network*

An important property of multi-agent simulations is that they can reveal emergent properties of social systems that cannot be predicted merely from analyzing individual behavior. The proper modeling of individual behavior to the level of sophistication and complexity required to achieve plausibility adds another level to emergence. Multi-scale emergence is a new property that may be observed by careful modeling of both intra-individual processes and interpersonal interactions. Although many connectionist models of intra-individual behavior have been proposed, very few connectionist models have appeared in the literature at multiple scales.

One exception is Talking Nets [4], which is a natural extension of the attitude model discussed in the previous section. Talking nets is a network of connectionist networks, where individuals can align attitudes by spreading activation through links connecting analogous nodes in neighboring individual networks. A critical parameter of the model is cognitive trust, which regulates the degree to which individual nets are able to influence each other. In a simplified view, cognitive trust increases as a function of the similarity of the evaluations between networks. Agents that disagree tend to filter each other's messages to a greater extent than if they tend to agree on issues. This is a simple but elegant solution to the problem of how to interconnect individual networks without dissolving individuality. The networks cannot be interlocked in such a manner that boundaries between individuals are not respected. One way to bias towards individuality is by setting a high threshold for cognitive trust to simulate the behavior of individuals that are resistant to social influence. On the other hand, cognitive trust can be extended liberally to members of a highly cohesive group in order to simulate the effects of collective action. At the opposite extreme, cognitive trust thresholds can be lowered to very low values, expressing a relative lack of individual cognitive boundaries, in order to simulate the interlocking effects of team cognition, such as one would expect to find in team sporting events or in military operations.

E. *Examples*

Figure 1 shows a fragment of a connectionist network which is used as a representation of an individual's attitude about a war against a foreign insurgency[†]. Direct links from the 'Fighting Insurgency' object to the positive (+) and negative (-) evaluation nodes represent the automatic or implicit activation of attitudinal response whenever the object is instantiated in semantic memory. Attributes such as 'winning' and 'increase security' are clustered around the attitude object. There are inward-bound links from the object to its attributes, and outward links connecting attributes to the appropriate evaluation node.

Since the model has a localist representation, all attributes are imparted fixed semantics that provide a concrete representation of the important constituent beliefs for evaluation in short term memory. Attributes can potentially receive inputs from external sources that simulate semantic priming of the concept. For example, news media reports that suggest that the war against the insurgency is being won would provide additional activation for the positive evaluation node; conversely, reports that the war is being lost boost the activation of the negative evaluation node.

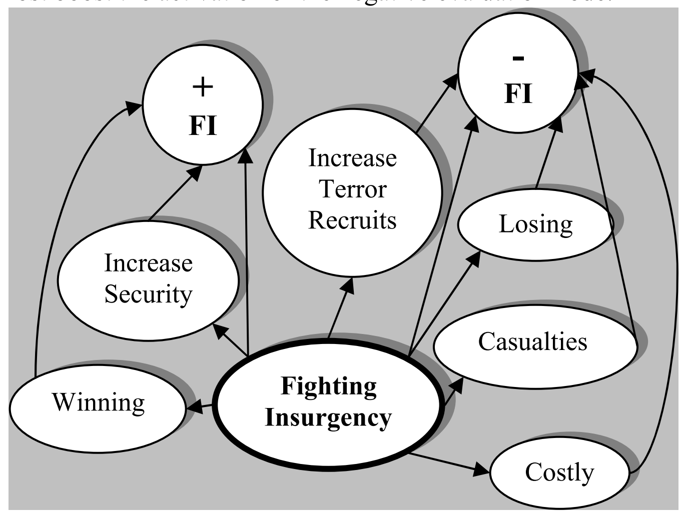

Figure 1: Individual attitude structure for fighting an insurgency. The attitude object is in bold type; attributes are shown in ovals with regular type; evaluation objects are depicted as circles.

In a more complex cognitive model, the number of objects and attributes towards/against those objects can be larger and inter-connected with more complexity. As an example, Figure 2 portrays an attitude structure consisting of four separate object nodes (fighting insurgency, president, economic recession and terror attack) together with their attributes and evaluation nodes. However, related attitudes within individuals are rarely independent. Rather, attitudes toward similar objects are usually highly correlated. This tendency reflects the property of attitudes toward consistency and coherence.

Goldstone and Janssen [17] stressed the importance of preserving interdependencies among cognitive elements as a partial justification to strive for greater sophistication of the internal representations of individuals in multi-agent simulations. We concur with their view, namely, "Concepts are not independent elements, but rather support and contextualize one another. Concepts gain their meaning by their relations to other concepts." Accordingly, we augmented the attitude network with additional links to model interdependencies among attitudes. The covariance structures that underlie attitude networks can be represented by same-level connections between attitude objects. These same-level connections, which go much beyond the kind of connections found in [3], trace the pathways of spreading activation due to priming mechanisms that operate on the basis of similarity.

[†] The graph structure is developed in the spirit of Van Overwalle and Siebler's framework. Other network representations are also possible, which could instantiate cognitive concepts in the form of nouns/gerunds in place of both objects and attributes. Due to space limitations here, we omit the alternatives in the connectionist network development methodology.

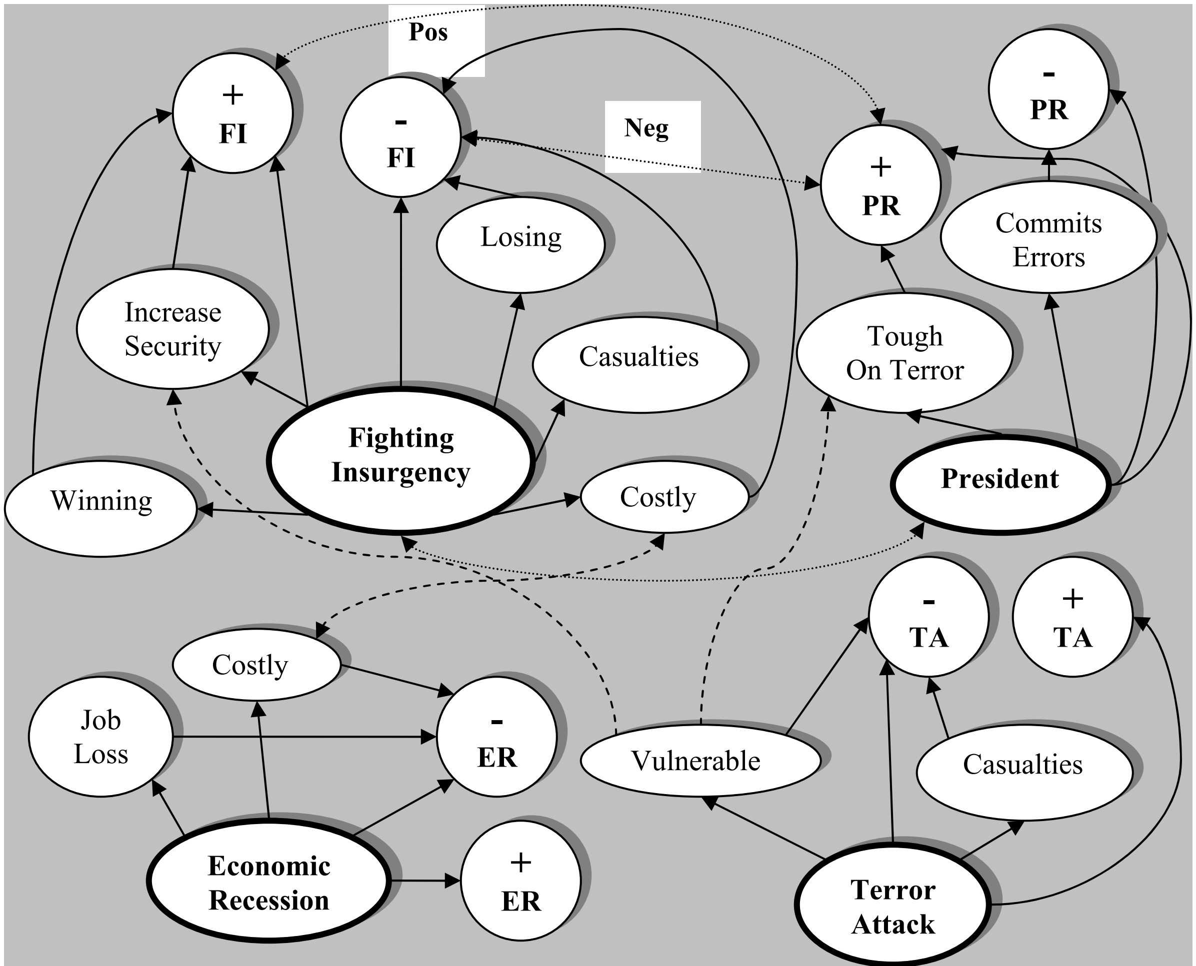

Figure 2: Elaborated connectionist attitude network for political opinions on national issues. Four attitude objects are shown. Object-object connections link 'President' and 'Fighting Insurgency'. Attribute-attribute connections link: (1) 'costly attributes for objects 'economic recession' and 'fighting insurgency'; (2) 'increase security' and 'vulnerable'; and (3) 'vulnerable' and 'tough on terror.' Between-object links are shown in dashed (attribute-attribute) or dotted (evaluation-evaluation and object-object) lines. Evaluation nodes are fully-interconnected in the model although only two links are shown to reduce clutter.

The phenomenon of "affect priming" is responsible for links between like-evaluation nodes for different attitude objects. Bower first suggested that cognitions and mood states may be linked in a single associative network. Therefore affect may selectively prime concepts to which it is connected. [18]. A positive evaluation of one object increases positive evaluations for other objects while negative evaluations are inhibited. The converse holds true for the effects of negative evaluations. This type of wiring may account for mood effects where all objects in STM receive identical coloration according to the prevailing mood.

Object-object connections reflect semantic similarity, where weight strength is positively correlated with similarity. Here we refer to semantic similarity in terms of associational strength. A semantic priming mechanism is responsible for triggering similar concepts following activation of an original concept. [13]. Referring to Figure 2, suppose that a news report about the war (that generates an automatic evaluation) reminds the individual of the country's president. Activation of the 'president' concept then spreads activity to an automatic evaluation. For one individual, a report that the country is winning the war against an insurgency primes the thought that the president is doing a great job. Another individual, upon hearing a news report indicating the war is being lost may recall that they disapprove of the president's job. Connections between attitude objects can produce unintended shifts of attention during attempts to persuade or cause a change in the attitudes of individuals.

Activation patterns in STM can also have side effects by actually changing short-term evaluations of other attitude objects. Again referring to Figure 2, a terror attack somewhere else in the world may remind an individual of the state of vulnerability of one's own nation. The attribute of 'vulnerable' semantically primes the presidential attribute 'tough on terror,' thereby increasing the activation of that node. As a consequence, the evaluation of the president is boosted and becomes even more positive. Simultaneously the increased activation or saliency of the 'vulnerable' attribute calls greater attention to the objective of the war against insurgency to 'increase security'. Therefore, more activation spreads to the positive evaluation of the war. Finally, the observation that an economic recession is costly could prime a similar observation about the war, thereby automatically reducing the level of support for continuation of the war.

IV. Applying Connectionism to Scale and Complexity

A. Methodological Solution

As illustrated in the examples, a single conceptual network based on well-conceived behavioral elements can capture a large variety of cognitive phenomena at once. In general, the initial model development can be split into two distinct parts: the first is the network development, and the second is the training and refinement of the network for different scenarios of interest. In the first, a concerted effort by subject matter experts can focus on the development of most pertinent networks in large classes of cognitive problems. The second part can be performed to either customize (train) the network to different sub-groups which can be as diverse as disparate countries, cultures or societies. The automatic accommodation of a variety of possibilities, afforded by a connectionist approach from the outset, provides the multitude of benefits in modeling large-scale scenarios.

Thus, coverage of heterogeneity is afforded not only within a group of interest, but also across groups. This solves three challenges: (1) richness of diverse behaviors *within* a large population can be realized as a consequence of the behavioral variety afforded by the connectionist network (2) hard-coding to a particular group such as a country or society is circumvented due to the greater reliance on the cognitive network, making it more plausible to reuse in another entirely different country or society, and (3) speed of modeling is achieved by reusing major portions of the basic cognitive modeling portions developed *a priori*. Additionally, automation is possible by employing training data on the network in an offline manner in cases where such training/historical data is available. Also, it is conceivable to more rapidly update the model if/when new real-time data becomes available to re-train the network in a rapidly evolving scenario.

B. Modeling Solution

The ability to support latent states and hidden linkages between cognitive elements makes the connectionist approach highly applicable to problems in which insights into certain behaviors are lacking in the literature. The modeling flexibility given by the network can be great when the set of cognitive concepts is relatively well understood but their inter-relationships are either poorly understood or are highly diverse across different subject populations.

More importantly, the support of hidden states and linkages makes the modeling framework less constrained by statistics-based inferences of surveys and polls while still being able to accommodate any additional new insights/data.

C. Data Sources for Connectionist Model Configuration

A major advantage of the connectionist approach is the potential to populate connectionist models with lower-level, more readily-available data, as opposed to higher-level behavioral data such as survey/poll data. Examples of fast data collection processes can be as diverse as satellite imaging, monetary flow data, Web search engine usage, and consumption data on resources such as food and fuel. Such proxy data could potentially be applied to quantitative assignation of arc weights in a connectionist model. For example, the quality of roads, housing and other infrastructure, can be inferred from imagery, using which the weights can be determined for the linkage of attitudes to attributes representing infrastructure improvements. Similarly, variations in non-cognitive data such as financial flows, or the levels of crop output, can be used to train the weights on current levels of satisfaction or dissatisfaction of subject populations. As another example, a period search of local newspaper articles can suggest the type and frequency of exposure of the local population to political messages linking specific attributes to objects. Representative links from non-domestic regions might be "America supports democratic movements" or "America cares only about political stability." Many other sources are available to help quantify specific attribute-object weights. The arc-weights can be further refined through adaptive learning processes by exposing individual connectionist networks to external stimuli representing political, economic and cultural influences that are available in the historical record.

V. Summary and Future Work

We identified and proposed the adaptation of a connectionist modeling-based methodology to deal with the new demands arising from recent interests in modeling behavioral elements such as attitude formation in large populations. While the initial amount of research effort in pursuing a connectionist modeling methodology is of a relatively high-risk nature, the potential payoff is great. The conventional reliance on subject matter experts for the development and refinement of rule sets could be reduced or eliminated. Such elimination can result in significant savings in time and financial resources. The sufficiency of rapidly obtained data (e.g., satellite images), and non-reliance on statistically-weak or high-latency data sources such as surveys/polls can translate to highly timely, nimble, and potentially more accurate, analyses and predictions.

The potential for large increase in speed of modeling, configuration, re-configuration, refinement and customization afforded by connectionism makes the connectionist paradigm one of the most important modeling approaches to be investigated to answer the critical behavioral questions being asked for actionable insights on aggregate populations of cognition-rich individual elements.

Future work would focus on design and implementation of connectionist agent models in three socio-cognitive domains that are critical to modeling social systems: attitude formation and change, group decision making, and leadership dynamics. Examples include simulation-based studies such as: (1) opinion change in regional populations following decisions by military commanders that affect socio-cultural aspects of community life; (2) decisions political organizations are likely to take under time pressure given a repertoire of actions and environment influences; and (3) emergent leader-follower behaviors in globally distributed groups or virtual communities that cohere primarily by electronic communication. Although prototype connectionist models exist, none is well-adapted to model these domains in their current state. For example, the cognitive trust model requires the addition of social diffusion processes to model the spread of information and influence. A significant technical challenge is in model validation, which can be envisioned to assume a mostly qualitative nature, focusing on the ability of models to account for robust phenomena described in the social psychology literature.

REFERENCES

[1] J. Epstein, "Modeling Civil Violence: An Agent-based Computational Approach," *PNAS,* vol. 99, pp. 7243-7250, 2002/05/14 2002.

[2] M. Page, "Connectionist Modelling in Psychology: A Localist Manifesto," *Behavioral and Brain Sciences,* vol. 23, pp. 443-512, 2000/01/01 2000.

[3] F. Van Overwalle and F. Siebler, "A Connectionist Model of Attitude Formation and Change," *Personality and Social Psychology Review,* vol. 9, pp. 231-274, 2005//01/01 2005.

[4] F. Van Overwalle and F. Heylighen, "Talking Nets: A Multiagent Connectionist Approach to Communication and Trust Between Individuals," *Psychological Review,* vol. 113, pp. 606-627, 2006/01/01 2006.

[5] Stanford-University, "Connectionism," in *Stanford Encyclopedia of Philosophy*. vol. 2009, 2009.

[6] C. Castelfranchi, "Through the Agents' Minds: Cognitive Mediators of Social Action," *Mind & Society,* vol. 1, pp. 109-140.

[7] R. Sun and I. Naveh, "A Cognitively-Based Simulation of Simple Organizations," in *Annual Conference of the Cognitive Science Society*, Stresa, Italy, 2005.

[8] W. Jager and M. A. Janssen, "The Need for and Development of Behaviorally Realistic Agents," *Lecture Notes in Artificial Intelligence (LANI 2581),* pp. 36-49, 2002.

[9] D. E. Rumelhart, G. E. Hinton, and J. L. McClelland, "A General Framework for Parallel Distributed Processing," in *Parallel Distributed Processing*. vol. 1 Cambridge, MA: MIT Press, 1987.

[10] E. R. Smith and F. R. Conrey, "Agent-based Modeling: A New Approach for Theory Building in Social Psychology," *Personality and Social Psychology Review,* vol. 11, pp. 87-104, 2007.

[11] E. R. Smith, "Distributed Connectionist Models in Social Psychology," *Social and Personality Psychology Compass,* vol. 3, pp. 64-76, 2009.

[12] D. Van Rooy, F. Van Overwalle, T. Vanhoomissen, C. Labiouse, and R. French, "A Recurrent Connectionist Model of Group Biases," *Psychological Review,* vol. 8, pp. 28-61, 2003.

[13] D. E. Meyer and R. W. Schvaneveldt, "Facilitation in Recognizing Pairs of Words: Evidence of a Dependence between Retrieval Operations," *Journal of Experimental Psychology,* vol. 90, pp. 227-234, 1971/01/01 1971.

[14] B. M. Monroe, "A Connectionist Model of Attitude Strength and Change," in *Annual Conference of the Cognitive Science Society*, Nashville, TN, USA, 2007, pp. 1313-1316.

[15] R. E. Petty and J. T. Cacioppo, "The Elaboration Likelihood Model of Persuasion," *Advances in Experimental Social Psychology,* vol. 19, pp. 123-205, 1986/01/01 1986.

[16] M. Fishbein and I. Aizen, *Belief, Attitude, Intention and Behavior: An Introduction to Theory and Research*: Addison Wesley, 1975.

[17] R. L. Goldstone and M. A. Janssen, "Computational Models of Collective Behavior," *Trends in Cognitive Sciences,* vol. 9, pp. 424-430, 2005/01/01 2005.

[18] G. H. Bower and J. P. Forgas, "Mood and Social Memory," in *Handbook of Affect and Social Cognition*, J. P. Forgas, Ed., 2001, pp. 95-120.