

NEXUS: AN INTELLIGENT AGENT MODEL OF SUPPORT BETWEEN SOCIAL GROUPS

D. DUONG, * US Office of the Secretary of Defense
R. MARLING, Marine Corps Combat Development Command
L. MURPHY, US Office of the Secretary of Defense
J. JOHNSON, US Office of the Secretary of Defense
M. OTTENBERG, US Office of the Secretary of Defense
B SHELDON, Marine Corps Combat Development Command
S STEPHENS, Marine Corps Combat Development Command

ABSTRACT

Nexus is an intelligent agent model focused on the support of social groups and organizations for each other and blame for departure from a social contract as evidenced by actions. It is a model of the attribution of blame for events based on trustworthiness. Agents representing leaders of groups look for breach of contract and keep track of a network of supporters. They choose whom to support depending on interpretations of past events. They perceive events depending on trust, reinterpreting past events in light of the present and visa versa. Nexus uses a model of the Boltzmann machine neural network, for the mind of each agent. It is based on interpretive social science, and the narrative paradigm in particular. Results are presented on a study of an insurgency using data collected from subject matter experts. This data includes support levels of ten relevant segments of population for each other, ideological similarity between groups, relevant historical events, and how groups might react towards the government if the US comes in to help during a natural disaster. The point of the simulation is to predict how disruptive the US aid would be if military personnel took an active role. The output data is the level of support for every group for every other group. It was found that, given historical events, a direct action by the US government only cause one group to like the insurgency a little more than they would have had the US government chosen an indirect approach to disaster relief..

Keywords: Interpretive Social Science, Narrative Paradigm, Boltzmann Machine, Constraint Satisfaction, Social Simulation, Neural Network, Irregular Warfare, Insurgency

* Corresponding author address: Deborah Duong, OSD/PA&E SAC1401 Wilson Blvd., Suite 300, Arlington, VA22209. email: Debbie.duong.ctr@osd.mil

INTRODUCTION

Nexus is an agent based model designed to simulate scenarios of irregular warfare. It is based on the narrative paradigm (Fisher), as agents look to relevant historical actions, current support networks, and ideological closeness to create a coherent view which calculates support levels for other agents. Agents try to make a story that is coherent with their historical context, and in doing so, may minimize apparent facts that don't make sense with the rest of the story, in accordance with cognitive dissonance theory (Festinger). Nexus is used for irregular warfare because the emphasis on contract keeping can be used for divisive strategies of nonviolent conflict. It takes into account higher orders of support, so that present enemies that are potential supporters, or present supporters that are potential enemies, may be identified.

THE MIND OF AGENTS

The Boltzmann Machine

The Boltzmann machine, a variety of the Constraint Satisfaction genre of neural networks, is used to represent the agent's minds, one for each agent. The Boltzmann machine is good at representing interpretations and reinterpretations of evidence. For example, the Boltzmann machine can model the interpretation of a Necker cube with a face as either in the front of the cube, or in the back of the cube, but as not both at the same time. It does this by making the belief that one vertice is in front evidence for believing or disbelieving that another vertice is in front (Simon). In Nexus, the Boltzmann machine is used to represent an evolving interpretation of evidence and blame, and its effect on levels of support. The paradigm shift, whether it be the shift that occurs when seeing a Necker cube in a different way, or the shift that occurs when facts are reinterpreted so that different parties are seen as responsible, is the same consonance seeking process. Constraint satisfaction networks have been used successfully to model how people see social situations (Duong and Reilly; Sallach; Thagard).

The basic agent of the Nexus model is the social group, which is a group of persons (whether they are organizations or not). Each social group has one Boltzmann machine that it uses to take all factors into account in its decision of whether it supports another social group.

The neurons in the minds of agents are of three types, as illustrated in figure 1.

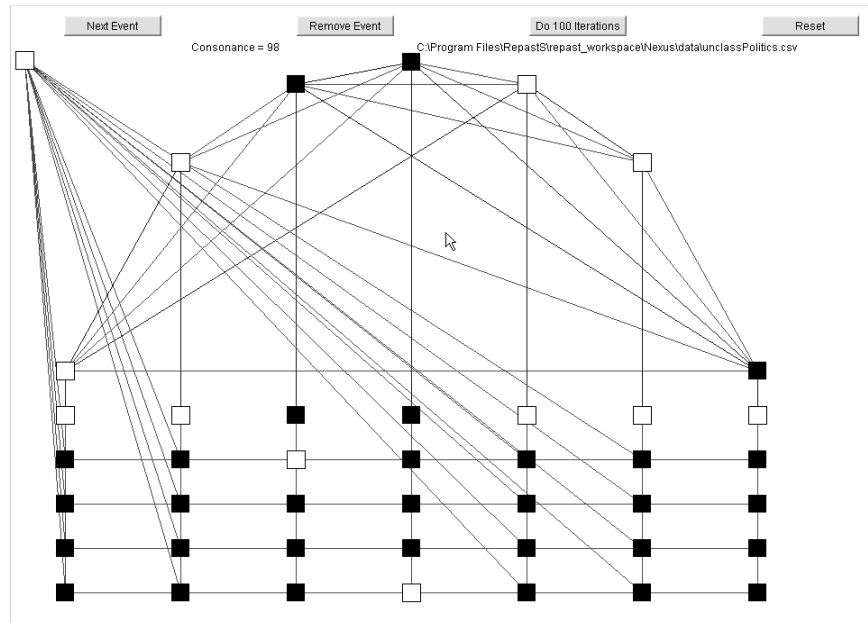


FIGURE 1 The Nexus GUI, of the neural mind of a single agent. The nodes along the top layer have been put into an arch so that their connections may be seen. Connections are red for inhibitory, and blue for excitatory. The columns represent the social groups. The top layer contains the support nodes. The second layer (which does not have connections between nodes) contains the trustworthiness nodes. Each layer below that is for a single historical event, and contains blame nodes corresponding to the amount of blame a group is given for an event. To the left is an input node, that holds objective evidence for the blame of each group for events before the “spin” the mind will place on it.

The Nodes of the Network

1. Support : An node for the level of support for each other social group.
2. Trustworthiness: A node for how much each social group is perceived as a keeper of social contracts.
3. Blame: A node for the belief that this social group performed a particular event, for every (social group X event).

The Architecture of the Network

1. Support nodes. These nodes output whether the social group owning the net supports another social group or not. These nodes have mutual excitation with the nodes of groups that publicly support each other and mutual inhibition with the nodes of groups that publicly lack support for each other. The weights change in the network depending on changing public declarations of level of support between groups. The support nodes ensure that the groups that a group is supporting are taken into account in its decision of who to support (for example, the friend of my

friend is my friend, the enemy of my enemy is my friend), and enables agents to perceive accountability as something that is shared with their network of support.

2. Trustworthiness nodes. The trustworthiness node for a social group has a mutual inhibition with all of the blame nodes for that particular social group. That means, if a group performs an adverse action, it is not generally perceived as trustworthy, and if it is trustworthy, it does not tend to be perceived as performing adverse actions. Trustworthiness nodes have a mutual excitation with support node of the social group, meaning that if a social group is trusted it tends to be supported, and if it is supported it tends to be trusted. There is also an input node to the trustworthiness nodes that corresponds to ideological similarity, so that having a set of beliefs and agreed upon practices is taken into account into the estimation of whether a group is a contract keeper.

3. Blame nodes. There are (social group X event) blame nodes, with sets of blame nodes for individual events appearing in rows and for individual social groups appearing in columns. There is a constant excitation applied to each of the blame nodes through an input node. The blame nodes are lit in proportion to the degree of hard evidence for the fault of each social group, before the spin that the mind puts on it. The total energy the blame nodes for an event are lit is in proportion to the severity of the event. The constraint that an action tends to be performed primarily by one entity is expressed by negative inhibition between the blame nodes within a row that represents a single event. Because of this constraint, if an event is blamed on one party, it tends to let another party off the hook. The Boltzmann machine can measure cognitive dissonance, or the spin that the mind places on an event, pulling it away from where contrary evidence, through the calculation of “goodness” or consonance of the net. That is, even though hard evidence supports blaming one group for an event, relations of support and trust may cause another group to be blamed.

Blame for one event is connected to blame for another event only indirectly, through the trustworthiness node. If a new event is determined to be a group’s fault, and the group’s trustworthiness falls, then evidence in the past for events can be reinterpreted and blamed on that group even if they were blamed on another group before.

Running the Simulation

When an action happens, it has some blame attached to it that represents objective evidence that an event was caused by a group, in a magnitude that reflects how much the event hurt or helped the group doing the thinking. The leader thinks, taking into account the whole picture of all the groups behaviors, their affinities towards other groups, and then makes a new public declaration of support. This declaration increments the support level mappings in each leader’s mind, preparing the leader for the next action when it will think and declare new support levels.

EMERGENT SOCIAL PHENOMENA

Nexus is a simulation based on first principles, and from which many types of tactics of irregular warfare may be modeled, including those discussed in Ackerman and Duvall’s book, *A Force More Powerful*. For example, Nexus can model the fact that a group has to worry about the upholding of an ideology with its peers. Gandhi’s revolution from India worked because Britain had to pay attention to its trustworthiness in keeping its ideology with international players. They had to worry about appearing hypocritical. Countries with ideologies that rationalize

violence, whose allies support the same ideas, such as Nazi Germany, may not be so afraid of their reputation, and crackdown on a protest. A group could think that its peers would not support it if it did not uphold the binding ideology as a social contract between itself and other players. In the case of India, knowledge of the ideological break with an innocent, non enemy power, would affect the reputation of the British in the model, forcing it to keep an ideologically correct social contract with India as well. In Nexus, every agent has a model of its perceptions of the support networks of other agents, and making knowledge of ideological breaks public through a non violent warfare campaign, affects support levels as agents worry about the keeping the trust of their allies.

As in irregular warfare tactics, the network may be manipulated to separate a regime from its supporters. For example, if the analyst inputs into Nexus an IO campaign of adverse events to be blamed on the police (such as “Rodney King” style videos) , then the regime may break their support of the police, to keep the support of the citizens, but at the same time causing services to not be delivered. To input this into the model, the blame nodes would be lit against a police group for the event. This may cause the regime to cut ties with the police, creating another adverse event that the police would blame the government for.

The support for groups of similar ideology and ethnicity would tend to affect trustworthiness of a group when they break contract with particular groups and ethnicities. It may affect those with similar ideology more than those less similar, as in social theories which stress the importance of empathy in the success of irregular warfare. If groups tend to judge other groups by the same standards, and have the same opinions, it causes them to have similar friends and enemies and tend towards mutual support. “Cognitive Liberation” may be simulated by a change in the ideology by which one judges how adverse an event is. For example, to simulate nationalism in a majority ethnic group for securities sake in the face of a possible civil war with a minority ethnic group, the collateral damage caused by a violent insurrection would be seen as an adverse event, but the minority ethnic groups might not blame the insurgent group that caused the action directly, as much as the government, on the basis of similar support networks as well as an accumulation of blame for the government on adverse events in the past. However, if the government was kinder to the majority ethnic group in the past, they may put a spin on the interpretation of events against the insurgents.

EXAMPLE RUN

Nexus has been run on classified scenarios to study questions of irregular warfare at the Office of the Secretary of Defense. It has also been run on unclassified but sensitive scenarios at the Marine Corps Combat Development Command. The object of the Marine Corps scenario is to predict the effect that the greater presence of the United States would have on the support for the government and the insurgents in a country, in case it had to help that country during a natural disaster. The US could take a direct approach to assistance, or an indirect approach. The groups include displaced persons, the urban poor, the urban middle class, old money, illicit organizations, the police, the army, the church, the government, and the insurgency. Data was obtained through two subject matter experts (SMEs) that knew the details of the culture and the particular group’s history in a single province. SMEs estimated support levels, what those levels might be in both direct and indirect cases, and ideological similarities of the groups. The SMEs also described 22 historical events that were important to the groups in determining their present feelings for each other. These historical events are from different time periods, as cultures can

consider something that happened thousands of years ago important. Historical events made up the cultural narrative and identity of the social groups. When Nexus was run, it was found that even though all groups except displaced persons were for the government and against the insurgents, the structure of their support for each other combined with historical events made the government somewhat vulnerable. All groups except the old money and the displaced persons changed their attitudes slightly more towards the insurgency when the US helped. It did not matter very much which kind of help. None of the groups, except for the urban middle class, had any different support levels for the government or the insurgents when direct and indirect action was compared. Only the urban middle class liked the insurgents a little bit more when the US action was direct than when it was indirect.

CONCLUSION

Nexus has been applied to real world scenarios of information operations and irregular warfare. It is one of the only tools that takes into account the historical consciousness of a people when explaining their actions. It is also unique in that it shows how new actions can influence a group to change their interpretations of the causes of their fortunes and misfortunes, and how these interpretations affect their alliances. Furthermore, it can reveal hidden vulnerabilities to changes in alliances due to higher orders of support levels and the entire historical picture of all parties. It has been suggested that Nexus will, in the future, be combined with Pythagoras for studies in stability operations.

REFERENCES

- Ackerman P., and J. Duvall, 2001, *A Force More Powerful: A Century of Nonviolent Conflict*, New York: Macmillan.
- Duong, D., and K. Reilly, 1995, "A System of IAC Neural Networks as the Basis for Self-Organization in a Sociological Dynamical System Simulation," *Behavioral Science* Vol. 40, number 4, pp. 275-303.
- Festinger, L., 1957, *A Theory of Cognitive Dissonance*, Stanford, CA: Stanford University Press.
- Fisher, W., 1984, "Narration as Human Communication Paradigm: The Case of Public Moral Argument," *Communication Monographs* 51, pp. 1-22.
- Sallach, D., 1988, "A Comparison of Parallel Architectures: Neural and Social Models of the Mind," *Proceedings of the Second ACM Symposium on Artificial Intelligence*, Norman, OK.
- Simon, Dennis. "The Boltzmann Machine Necker Cube Example," <http://www.cs.cf.ac.uk/Dave/JAVA/boltzman/Necker.html>.
- Thagard, P., 1999, *Coherence in Thought and Action*, Cambridge, MA: MIT Press.