

A Semantic Differential Approach to Incorporating Qualitative Data into Nexus, an Interpretive Agent Model of Support Between Social Groups

Robin Marling¹, Deborah Duong², Bob Sheldon¹, Steve Stephens¹,
Lauren Murphy², Jeffery Johnson², Michael Ottenberg²

¹ MCCDC/OAD Quantico, Virginia 22134

²OSD/PAE Simulation Analysis Center, 1225 South Clark Street Suite 300,
Arlington, Virginia 22202

robin.marling@usmc.mil, debbie.duong.ctr@osd.mil

Abstract. The Semantic Differential technique is applied to the words in subject matter expert descriptions of the feelings of one social group for another in a city in Colombia. Evaluation, Potential and Action scores are converted to Support Levels between groups and put, along with historical events that explain these levels, into an agent based model. Each agent in the model represents a group, and calculates its support level in a new situation taking account the network of support levels, similarity, and historical relationships between groups. The calculation is performed with a constraint satisfaction neural network.

Keywords: Semantic Differential, Attitude Measurement, Verbal Method, Agent-Based Simulation, Neural Network, Boltzmann Machine, Support Levels, Irregular Warfare

1 Introduction

This paper presents study of an Irregular Warfare scenario in Colombia, performed under the collaboration of the Office of the Secretary of Defense and the Marine Corps Combat Development Command. It describes the application of Osgood's Semantic Differential technique for deriving numerical scores of support levels between population groups from verbal descriptions, to input into a simulation of irregular warfare. It also describes an agent based simulation that uses these scores to predict future support levels. The Marine Corps Combat Development Command Operations Analysis Division (MCCDC/OAD) performed the data collection, and the Office of the Secretary of Defense Program on Analysis and Evaluation (OSD/PAE) developed the agent based simulation that used the data.

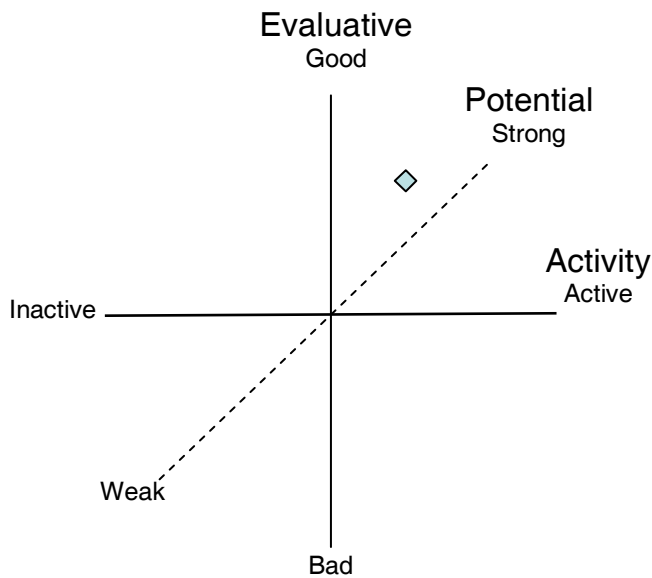
2 Data Preparation

The data is focused on the city of Buenaventura, Columbia which is the crossroads of both a violent insurgency and the blossoming illicit drug trade. With the assistance of subject matter experts (SMEs), MCCDC/OAD divided Buenaventura into eight population segments, or population groups, to be modeled. The SMEs included two historians, one political scientist and one intelligence specialist, all of whom have extensive time living in Colombia.

Semantic Differential theory was used to capture the affinity of the eight population groups for each other (Osgood et al, 1957). Semantic Differential theory posits that human beings make judgments based on three factors: Evaluative, Potential and Activity (EPA). Evaluative is a value judgment: is object A good or bad? Potential is a potency judgment: is object A weak or strong? Activity is an action judgment: is object A active or inactive? For example, a bunny rabbit would be good, weak and active; a bad boss would be bad, strong and active; and a parked car would be neutral, strong and inactive. Dr. David Heise quantified these values by surveying some 1000 respondents for numerical scores of about two thousand words, placing each E, P and A value on a Likert scale of -5 to 5.

To understand this theory, it is easiest to visualize words as data points in a three-dimensional space, where the Evaluative, Potential and Activity values represent (x,y,z). The good-to-bad axis represents the Evaluative factor, the powerful-to-powerless axis represents the Potential factor and the fast-to-slow axis represents the Activity factor.

Figure 1. Word value on a three dimensional scale



For our project, we wanted to devise a single numerical value to represent all three factors. We could have measured the length of the vector from (0,0,0) to (e,p,a) as:

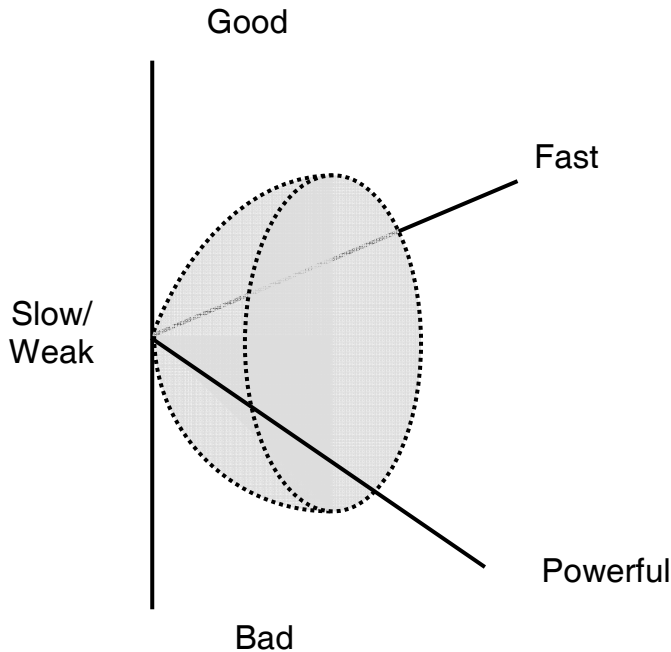
$$\sqrt{(Evaluative^2 + Potential^2 + Activity^2)}$$

However, that would give more a higher numerical value to a person/action thought to be bad, powerless and slow, than to a person-action described as good, moderately powerful and moderate activity, due to the -5 to 5 Likert scale. To avoid this, we decided to transform the Potential and Activity data so that it scaled from 0 to 10 and defined an input labeled Influence.

$$Influence = Evaluation * \sqrt{(Potential^2 + Activity^2)}$$

The results of this equation fall along a parabola similar to this:

Figure 2. Word value using Influence equation



Using this methodology, the following table was created for a list of modifiers. The first column is the original Evaluation score, the second and third columns are the modified Potential and Activity scores and the fourth column is the applied Influence equation:

Figure 3. Influence (Support Level) Scores for individual words

	Evaluation	Potential	Activity	Influence
Abusive	-2.62	5.26	5.76	-20.44
Unreliable	-2.17	3.54	5.40	-14.01
Hopeless	-2.03	3.08	4.29	-10.72
Boring	-1.73	3.53	3.61	-8.73
Dull	-1.42	3.63	3.38	-7.04
Rich	1.24	7.12	5.12	10.87
Powerful	1.18	7.46	6.07	11.35
Wealthy	1.47	6.93	5.11	12.66
Ambitious	1.67	6.79	6.62	15.84
Competent	2.10	6.71	5.51	18.23
Talented	2.22	6.77	6.06	20.17
Friendly	2.69	6.62	6.02	24.07

So, a population group would be repelled by another population group they thought of as abusive and unreliable; uninterested in a group described as dull; and attracted to a group they feel is talented and friendly. Group A having negative influence on Group B would indicate that Group A is repulsive to Group B, therefore Group B would reject Group A in their actions and decisions. If Group A had null (close to zero) influence, it would mean that Group A has little influence on Group B and

isn't considered in their actions or decisions. Group A having positive influence means that Group A is attractive to Group B and that Group A's goals and desires are taken into consideration in Group B's actions and decisions.

An interesting example of this is the difference in public reactions to controversial talk show hosts vs. controversial speakers on street corners. While both speakers are equally controversial, the Potential of the talk show host is greater and therefore reactions are more dramatic towards him or her.

2.1 Data Technique applied to a Colombian Town

The major benefit to using the semantic Differential is that Human population SMEs are comfortable with providing input using this method. We can avoid questions that require SMEs to directly assign numbers to various groups and instead ask questions such as: "how does Group A describe Group B?", and "how would Group A react to Event C?" By following the SME's conversation, we can pick out the salient words or phrases that they use in explanation and convert them to numeric values that can be used for modeling purposes.

In the Colombia scenario developed by the MCCDC/OAD IW study, two of the identified population groups were named Old Money and Urban Middle Class. When SMES were asked how the Old Money would describe the Urban Middle class, they used the following words: active, demanding, knowledgeable, outspoken and poised.

Figure 4. Influence (support level) scores for Old Money's view of Urban Middle Class

	Evaluation	Potential	Activity	Influence
active	2.07	6.53	7.28	20.25
demanding	-0.59	5.66	5.79	-4.74
knowledgeable	2.21	6.92	5.02	18.90
outspoken	0.57	6.31	6.30	5.05
poised	1.35	6.16	5.01	10.72

From this list, we can surmise that Old Money sees the Urban Middle Class as mostly good, with moderate potential and moderate activity. We can assume that the opinions and actions of the Urban Middle Class will influence the Old Money to a greater degree than the actions and opinions of several other groups.

2.2 Limitations of the Data Technique

There are, of course, limitations to this method. Major detractors from the Semantic Differential subscribe to theory called Nominalism, who posit that everyone defines words individually and that there is no universal definition of a word that could be quantified. Surveys conducted by social scientists have shown that words have different semantics from country to country, but within certain national boundaries and time frames, had pretty solid signal to noise ratios. Individual interpretation was also mitigated in this study by only using American citizens who had studied in the United States as SMEs and only using data from universities in the United States, which appeared fairly consistent. The data used to define Evaluative, Potential and

Activity values in the IW Study was derived from surveys of college students, not cultural experts, so we also made the assumption that the difference between the two groups (college students and human population SMEs) was small enough to warrant using the database. Our formula for Influence was created for modeling convenience. However, similar approaches have been used by other social scientists with great success. The database of 516 modifiers and over 1000 verbs that we used was large, but not all inclusive. MCCDC/OAD attempted to mitigate this by providing immediate visual feedback to the SMEs, to vet the input they were giving.

The most profound drawback for using this technique is that it does not provide the context in which answers are given. As one SME told us, “When I am picking words, I am giving you knowledge (data), but when I explain my answers, I am giving you wisdom (context).” The lack of context is mitigated in the input to the Nexus agent program with the inclusion of an additional data set on the historical events behind the influence scores.

Still, we feel that this method has allowed us to have confidence in the data we have collected. Previous methods of collecting qualitative data for modeling efforts forced SMEs to contort knowledge into uncomfortable formats, such as using a numerical scale, selecting from a list of narrowly defined options, or forcing group consensus. By using Semantic Differential theory, we have solicited data from SMEs using their words and their ideas. Hopefully, this is a step in the right direction of incorporating the qualitative into the quantitative. This effort, we feel, is worth further investigation.

3 Nexus, the Agent Based Model


Nexus is an agent based model designed to simulate scenarios of irregular warfare. It 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. Nexus models 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 is used for irregular warfare because the emphasis on contract keeping can be used for divisive strategies of nonviolent conflict, and scenarios in which identity politics is important.

Nexus uses the Boltzmann machine neural network for the mind of each agent. It is based on interpretive social science, and the narrative paradigm in particular. In simulating the narrative paradigm (Fisher), 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 also 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.

3.1 The Mind of Agents: The Boltzmann Machine

The Boltzmann machine, from the Constraint Satisfaction genre of neural networks, is used to represent the agent's mind, one for each agent. The Boltzmann machine is good at representing interpretations and reinterpretations of evidence. A good first step to understanding what Nexus does is to understand the use of the Boltzmann machine on a simple example of cognition, the interpretation of the Necker cube, because the process of visual cognition is similar to the process of social cognition, but is easier to understand. The example below is from Dennis Simon's interactive "Boltzmann Machine Necker cube example" at <http://www.cs.cf.ac.uk/Dave/JAVA/boltzman/Necker.html>, recommended to the reader for a more interactive experience. In this example, the Boltzmann machine models the way a human mind "switches" its interpretation of a Necker cube, with a face as either in the front of the cube, or in the back of the cube, but not both at the same time. See figure 5 for a picture of the Necker cube, to note how our mind seeks a consistent interpretation of evidence by switching back and forth between symmetrically consistent interpretations. The mind switches interpretation often because the amount of evidence for one interpretation is the same as the amount of evidence for another interpretation: if the visual interpretation (or the social interpretation) was less symmetrical, it would not "switch" as often.


Figure 5. The Necker Cube



The neural network that simulates the creation of a coherent picture out of the Necker cube comes with two subnets: an upper subnet that represents the idea that the front of the cube is facing downwards, and the lower subnet that represents the idea that the cube is facing upwards. The left side of figure 6 shows the two subnets without links, and the right side shows the net with the links. Each of the nodes represents evidence about a state of the world: in this case, a vertex's position on a cube. The colored links on the right side of the diagram represent the amount of support of one state for another, blue if its positive (excitatory) and red if its negative (inhibitory). For example, if we believe that the bottom left vertex is in the front of the cube, we have reason to believe that the bottom right vertex is also in the front of the cube, and so there is a blue link between the bottom left and bottom right vertices. On the other hand, since the top subnet says the bottom left vertex is in the front, and the bottom subnet says the bottom left vertex is in the back, and they can

not be both at the same time, there is a negative link between the two corresponding vertices on each subnet.

Figure 6. A Boltzmann Machine that simulates the mental act of interpreting of the Necker Cube. The left side is the network without the connections shown, and the right side is the network with the connections shown.



In order to simulate an interpretation, nodes are at first turned on randomly. In the example, nodes that are white are turned on, meaning there is evidence for the state, and nodes that are black are turned off, meaning that state is not believed. If a node is on, it tends to turn on those nodes it is linked to positively, and turn off those nodes that it is linked to negatively. The pictures depict a random first state, but after many iterations of nodes turning other nodes on and off according to evidence, the net chooses either the upper or the lower subnet. Either the upper net becomes completely lit, and the lower net completely off (meaning the mind is thinking that the front of the cube is facing downwards) or the upper net turns completely off, and the lower net becomes completely lit (meaning the mind is thinking that the front of the cube is facing upwards). The net settles down on a single thought, which can change as soon as another interpretation is made. In Nexus, the net settles down on a single thought of social support.

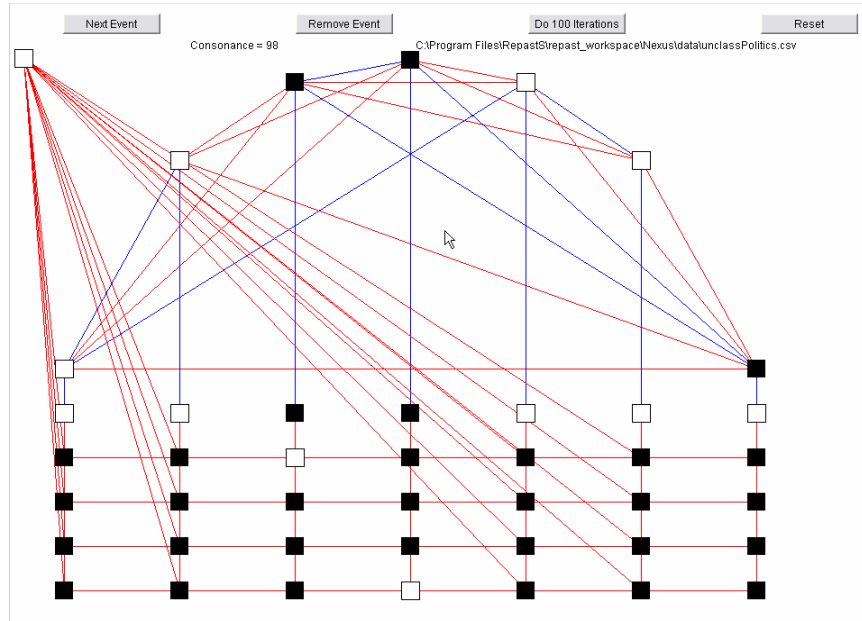
3.2 The Boltzmann Machine in Nexus

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 social groups are seen as responsible, is the same consonance seeking process. Constraint satisfaction networks like the Boltzmann machine have been used successfully to model how people see social situations (Duong and Reilly; Read and Miller; 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 7.

Figure 7. 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. The support nodes are in accordance with Heider’s balance theory, which states that

people experience cognitive dissonance when they are the friend of their enemy's friend, or the enemy of their friend's friend.

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 and their affinities towards other groups. As the leader thinks, it may have to minimize certain inconsistent evidence for historical events, or minimize the fact that it is similar to a group that is its enemy, or minimize a friendship when circumstances force the group to be the enemy of its friend's friend. These less than ideal situations are made up for with rationalizations, as in Festinger's theory of cognitive dissonance. The leader then makes a new public declaration of support, based on its new thinking. 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.

3.3 Nexus as a Cognitive simulation of Irregular Warfare

Nexus simulates identity politics and the narrative paradigm in that its agents seek narrative coherence. According to Fisher, people think with narrative coherence, that is, they think in terms of stories in which the characters have consistent

behaviors and act according to an identity. Cognitive dissonance theory has a lot to do with identity coherence: according to Festinger, people tend to minimize the evidence for facts that go against a consistent picture of their identity: for example, smokers will often minimize the evidence that smoking is bad for their health. Both the narrative paradigm and cognitive dissonance theory play a role in identity politics, for example, the extremist view that Bush caused 9-11: extremists think so because they believe he is the kind of person who would do such a thing, and it goes with their narrative. They experience cognitive dissonance when they minimize the evidence that he did not do so. When Nexus puts together a picture of historical relations, it overlays consistency on history, so that blame for past actions is reinterpreted based on a coherent story.

Because it closely examines the cognitive reasons for support, Nexus is a simulation based on first principles 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.

3.4 Run on the Colombian Data

In addition to the MCCDC Columbia scenario, Nexus has been run on classified scenarios to study questions of irregular warfare at the Office of the Secretary of Defense. The object of the MCCDC scenario is to plan US humanitarian assistance response to a hypothetical natural disaster in Beuna Ventura. The US could take a direct approach to assistance, placing military assets on shore, or an indirect approach, placing assets on sea. The groups modeled 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.

In addition to the Semantic Differential data, the SMEs also described 22 historical events that were important to the groups in determining their present feelings for each other. These historical events, important to the cultural/historical consciousness of the 8 groups, are slavery, the assassination of Jorge Gaitan , the military dictatorship of Gustavo Rojas Pinilla, La Violencia, FARC terrorism, the US overthrow of Manuel Noriega's government in Nicaragua,, the Cimmaron movement, militia terrorism, the problem of the Land (poorly documented land ownership that made legal recourse to property violations difficult), Failure of 1984 Truce and systematic assassinations of UP officials, the indigenous persons movement, the 1991 constitution, the 1993 amendments to the constitution, the expansion of the drug trade, the dissolution of the Cali Cartel, the 1997 Police Reform, the 1998 Economic Slump, the failure of the 2000 FARC peace agreement, Plan Colombia, the 2002 Military Reform, the 2005 Disolution of the AUC and the Rise of Bacrim

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.

3.5 Nexus 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. Nexus is hopeful in that it examines the way that new actions can cause a paradigm shift in the way one group sees another, despite historical enmities, dissimilarities, and alliance structures.

4 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.

Heise, D. R. 1970. The semantic differential and attitude research. G. Summers (Ed.), *Attitude Measurement* Chicago: Rand McNally, pp. 235-53.

Osgood, C.E., Suci, G., & Tannenbaum, P. (1957) *The measurement of meaning*. Urbana, IL: University of Illinois Press

Read, Stephen and Lynn Miller. "On the Dynamic Construction of Meaning: An Interactive Activation and Competition Model of Social Perception," *Connectionist Models of Social Reasoning and Social Behavior*. London: LEA 1998.

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.