OCCAM: Ontology-based Computational Contextual Analysis and Modeling

Srini Narayanan[1], Katie Sievers[2], Steve Maiorano[2]

[1] International Computer Science Institute (ICSI) and University of California, Berkeley
1947 Center Street, Berkeley, California 94704.
snarayan@icsi.berkeley.edu

Abstract. The ability to model cognitive agents depends crucially on being able to encode and infer with contextual information at many levels (such as situational, psychological, social, organizational, political levels). We present initial results from a novel computational framework, Coordinated Probabilistic Relational Models (CPRM), that can potentially model the combined impact of multiple contextual information sources for analysis and prediction.

Keywords: Bayesian modeling, cognitive models, agent models, computational sociology

1 Introduction

The ability to capture contextual influences formally (where context includes social structures, practices, and norms) is of great importance in today's globalized and multi-cultural environment, not only to foster intercultural understanding but also to provide a collaborative and computationally tractable way to predict potential conflicts and crises. Current modeling tools are primarily based on mining patterns from data and ignore important social, psychological, and cultural dimensions that often determine cause, motivation, and intent. Knowledge pertaining to cultural factors is informally (if at all) applied based on an individual's instinct and is neither indicative of corporate experience nor shared. There is thus a pressing need to enhance the current analysis and policy-making framework with tools that operationalize the acquisition and use of the best cultural and social scientific data and make available this knowledge to a collaborative process.

This paper reports on initial steps and preliminary results addressing this state of affairs. Specifically we address the following issues.

- How do we systematize the encoding and application of psychological, social/cultural factors, norms, and practices to analysis and prediction?
- Can we formally model the impact of context for analysis? Our focus is on formal operational models that support analysis, explanation, and prediction.

[1] Srini Narayanan is the contact author.
[2] Katie Sievers and Steve Maiorano are affiliated with the Advanced Behavioral Sciences Group (ABSG) of the United States Government.

- Can we apply our methodology to support cross-cultural comparisons? We are exploring formal techniques that can shed light on highly sensitive factors, and important cultural similarities and differences that condition possible outcomes and responses of different groups in specific situations.
- Of central concern in today's world is the rapid operationalization of knowledge so we can bring it to bear on cross-cultural analysis. Here we are exploring ontology-based tools for knowledge entry as well as use advanced semantic extraction technology to mine relevant social and cultural information from different media sources including open media.

The rest of the paper is organized as follows. First, we state the background assumptions, findings, and research that guide our model building and analysis efforts. This is followed by a description of the computational modeling and simulation framework. We report on preliminary results from the application of the computational tools to two case studies involving motivational and cultural context for analysis and prediction. The first of the two models is based on an extensive threat assessment framework developed at the Monterey Institute for International Studies (MIIS). The goal of this model is to automate via a computer simulation a factor analysis protocol where the outcome variable is the likelihood of specific groups attacking critical infrastructure targets. A second model predicts the impact of specific international policies on coca-eradication efforts in the context of two similar countries, Bolivia and Peru. Of specific interest here was to identify similarities and differences in predicted response between the two countries and compare the model predictions to actual responses of the two countries.

2 Background

The early view of sociologists that cultural context was a "seamless web" [3], unitary and internally coherent across groups and situations has given way to the realization that depicts culture as fragmented across groups and inconsistent across its manifestations. The current view of culture is as complex structures that constitute resources that can be put to strategic use. Once we acknowledge that culture is inconsistent, it becomes crucial to identify units of cultural analysis and to focus attention upon the relations among them. Similarly, once we acknowledge that people behave as if they use culture strategically, it follows that the cultures into which people are socialized leave much opportunity for choice and variation. Thus our attention turns to ways in which differing cultural frames or understandings may be contextually cued.

Research from a variety of disciplines [2, 3, 7, 8, 9, 10, 11, 14, 15, 16] has focused on study and formalization of psychological and cultural framings and their situational evocation[3]. A great deal of convergent results from Cognitive Science, Sociology, Anthropology, and Neuroscience indicates that complex concepts are not categorical (cannot be defined by

[3] For more information on linguistic and cultural frames see http://framenet.icsi.berkeley.edu. For neural underpinnings and computational models see http://www.icsi.berkeley.edu/NTL

necessary and sufficient conditions).[4] Instead complex concepts exhibit a radial structure often with a prototypical member and a number of mappings and extensions [14, 15, 16]. Prototypes of categories could arise from various considerations including a) being a central category (others relate to it; amble and swagger relate to the prototype walk), b) being an essential feature that meets a folk theory (birds have feathers, lay eggs), c) being a typical case (sparrow is a typical bird), d) being an ideal positive social standard ("parent) or an anti-ideal negative social standard ("terrorist"), e) a stereotype (set of assumed attributes as in dumb blonde) or f) a salient exemplar (second world war as a just war).

Abstract concepts (such as democracy, freedom, love, mathematical concepts) are often metaphorically mapped from more experiential domains such as *force, spatial motion, and social cognition.* There is a lot of cross-cultural metaphoric knowledge that relies on distinctions in the experiential socio-cultural groundings of abstract and contested concepts [9, 10].

Over the last decade, we have been building models of conceptual acquisition and use which exploit these findings [4, 11, 12] on the structure of complex conceptual categories. Such models have been formalized and used to encode complex metaphors, frames and grammatical constructions [7, 8, 9]. We now describe the use of these models to provide a realistic framework for encoding complex motivational, and socio-cultural knowledge for use in analysis and prediction.

From the computational modeling perspective, we situate our efforts in the much larger matrix of tools and techniques that have or potentially can be applied to modeling context in all its forms. **Table 1** shows the various modeling frameworks and their capability to capture contextual background (including social and cultural) knowledge. Previous attempts to incorporate social and cultural knowledge include ad-hoc approaches including individual intuition, dynamic system theory, cellular automata or artificial evolution techniques, rational actor theories, game theoretic techniques, Bayesian analysis, and logical techniques. Each of these approaches addresses some aspects concerning the modeling and use of socio-cultural knowledge. For instance, game theoretic agents are able to model fairly complex social agents capable of sequential and iterative decision making in the presence of other social actors in an uncertain environment. However, the rational actor model which provides the theoretical underpinnings, has questionable ability in terms of modeling actual human motivation and intent which has led to a variety of divergences between theoretical predictions and actual data. Cellular automata and evolutionary techniques (including artificial life) capture the essential property of *emergence* in complex systems, where collaborative and competitive behavior between multiple agents (usually multiple generations of agents) leads to hierarchies, coalitions, segregations and other complex arrangements manifest in human groups. However, these approaches are notoriously hard to analyze and have yet to take into account the structural complexity of cognitive agents.

[4] An illustrative example is the concept *mother* with possible members including birth mother, biological mother, genetic mother, nurturant mother, surrogate mother, stepmother, etc.

Hence there is a need to unify these approaches and synthesize a new technique that is able to account for and utilize the observations that a) social knowledge is not categorical and requires the use of prototypes and radial categories, b) human perception, motivation, and social practices are fundamental, not epiphenomenal, and c) events and actors change, evolve, and adapt dynamically.

Table 1: Shown above are the various approaches used to model complex scenarios that include emergent behavior, dynamic situations with imprecise and incomplete information, where multiple social actors engage in decision making and coalition formation in the pursuit of shared interests and goal. Graphical models here include Bayes Nets, DBN based models as well as undirected models including Markov Random Fields (MRF) and Conditional Random Field Models (CRF). They also include standard neural networks trained using back-propagation techniques (Multi-layer Perceptrons) and recursive networks (SRN). Cellular Automata models include swarm intelligence models as well as spatial density function models. Hybrid CPRM are extended CPRM which incorpate both discrete and continuous state.

Technique	**Basic Technology**	**Complexity of Social Actors**	**Learning Adaptation**	**Continuous State**	**Dealing with incomplete information**
Non-Linear Dynamic Systems	**Differential Equations**	**low**	**low**	**yes**	**no**
Stochastic Processes	**Markov/Queuing models**	**Low/ Medium**	**Low-high**	**partial**	**Medium**
Graphical Models	**Bayesian Networks, MRF, CRF**	**medium**	**high**	**partial**	**medium**
Cellular Automata	**Evolution rules for agents**	**low**	**high**	**no**	**low**
Social Networks	**Graph analysis (small world)**	**medium**	**med**	**no**	**medium**
Game Theory	**Rational actor model**	**high**	**low**	**no**	**medium**
Multi-Agent systems	**Distributed AI/swarms/ Artificial life**	**high**	**low**	**no**	**medium**
Qualitative Techniques	**Naïve Physics**	**medium**	**low**	**partial**	**high**
Hybrid Automata	**Hybrid-system theory**	**medium**	**low**	**yes**	**high**
ICSI Hybrid CPRM	**Bayesian Hybrid-system theory**	**high**	**med**	**yes**	**high**

For several years we been developing a novel computational framework called Coordinated Probabilistic Relational Models (CPRM). The CPRM modeling framework is unique in being explicitly designed to model complex psychological, social, and cognitive context and the software program implementing the framework is capable of constructing a faithful, robust, flexible, interactive, and graphical computer simulation of the entire decision process. The framework is faithful in that it accurately models the process threads, event models and decision processes in the analysis protocols; robust in the sense of doing so over a wide range of parameter settings; flexible in being able to generate the best hypothesis consistent with partial, incomplete, and mutable data; interactive in that the system operation and parameter settings can be interactively changed while the software program is executing, and hypothetical ``what-if'' simulations performed; and graphical in that the model is a formal graphical structure that supports visualization of the decision process as well as exact quantitative analysis.

3 Modeling Approach

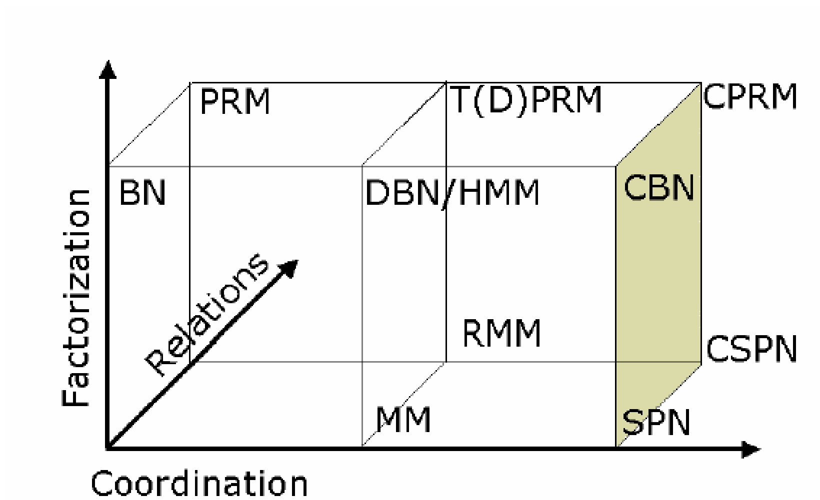

Fig. 1. Probabilistic Inference: The space of models. The origin corresponds to pure statistical models that are based on attribute value vectors (such as vector space models). The x-axis represents increasing sophistication is handling dynamics (from sequences to fully branching dynamics). The y-axis represents graph properties of the model and the use of independencies to make inference tractable (as in graphical models). The z-axis (into the plane) represents modelling expressiveness moving from propositional to fully relational representations.

Modern inference systems deal with the ambiguity and uncertainty inherent in any large, real-word domain using probabilistic reasoning. Such models have many advantages, including the ability to deal with missing and uncertain data. Bayesian networks have worked extremely well in moderate sized cases, but do not scale to situations of the size and complexity needed here to model analyze and infer aspects of complex systems. In general, reasoning about complex event structure, ambiguity, and dynamics requires modeling coordinated temporal processes and complex, structured states. A significant amount of work has gone into different aspects of overall problem.

Figure 1 maps out the space of relevant probabilistic modeling and inference techniques along three basic dimensions (extended from the description in (Anderson *et al.* 2002) [1]). The dimension along the x-axis (left-right) depicts the increasing dynamic component. In Figure 1, the x-axis corresponds to the complexity of the dynamics modeled, the y-axis (vertical going up) to the complexity of the state representation and the z-axis (into the plane) the richness of the relations. The origin of the space is an unstructured probabilistic state vector representation with no explicit temporal or relational information. Moving to the right along the x-axis, we get to linear temporal models of sequences. Markov Models (MM) are the most widely used technique to model such simple sequential processes. However, Markov models are fairly inflexible and representationally inadequate as models of actions. Moving further right, we arrive at a set of well developed graphical modeling approaches (such as Stochastic Petri Nets (SPN)) designed to model distributed dynamic systems with complex coordination, concurrency and resource constraints.

Moving from the origin up along the y-axis, we have Factor Models, Markov Random Fields (MRF) (the undirected version) and Bayes Nets (BN) (the directed version). Temporally extended Bayes Nets (TBNs, also called DBNs) model each time step in a sequence as a BN and links between state variables at different time steps capture the temporal dependencies between variables. Moving rightward, Coordinated Bayes Nets (CBNs) [12] combine the expressive action modeling framework provided by the SPN or CSPN (SPN with objects) based representation with the ability to model complex states provided by the BN framework. This model of action and its use in reasoning about actions was illustrated in [11, 12]. PRMs [16] extend the Bayes Net formalism to allow specification of a probability distribution over a set of relational interpretations.

We have been developing a novel computational framework called Coordinated Probabilistic Relational Models (CPRM) which result from a complete integration of PRM and our extended Hybrid Petri nets (HySPN), and are suitable for building and learning formal models of complex systems. The CPRM modeling framework comes out of a larger project at Berkeley that has been building neurally plausible computational models of cognitive and social phenomena (such language, motivation, intent) for over two decades (http://www.icsi.berkeley.edu/NTL). As far as we know, CPRM models are unique in their ability to model, learn and reason evidentially about dynamically evolving events in complex, uncertain environments.

Currently, an initial implementation of the framework (described in [11] and henceforth referred to as the ICSI simulator) has demonstrated the capability to go beyond finding generic and often trivial patterns in data to discerning correlations and activity that reflects purpose and intention. The simulator is the reasoning component of a question-answering system, AQUINAS (Answering Questions Using INference And Semantics) [13], designed to provide answers to complex questions requiring justification, causal inference, and hypothetical reasoning. An essential requirement for systems that attempt complex causal inference about purpose and intent is the ability to model context specific information in an uncertain, partially observable environment.

This paper reports on our first attempts to use the ICSI simulator to formally model the various process threads and the contextual influence of multiple uncertain systemic, social, organizational, and psychological factors. Analysts can interactively manipulate and conduct “what-if” simulations to assess the sensitivity of individual factors or combinations of factors on the likelihood of specific threat hypotheses and attacks.

4 Initial Results

We report on two case studies relating to the use of the model. The first represents work performed in conjunction with the Monterey Institute for International Studies (MIIS) and is based on making operational in a computer model an elaborate threat assessment framework developed at MIIS. The goal here was to replicate in a computer model the decision making process that individual analysts engaged in to evaluate the threat potential of specific cases involving specific groups. The second model was an attempt to capture complex social and cultural factors, combine their contextual influence with other physical, systemic, and motivational information sources supporting different hypotheses in a probabilistic framework. The goal here was to be able to explore the explanatory power of these different information sources in predicting similarities and differences in the response of different nation-states (in this case Bolivia and Peru) to specific international policy objectives (in this case coca-eradication).

4.1 An operational model of motivation and intent for threat assessment

One application resulted in the first operational model of the Determinants Effecting Critical Infrastructure Decisions (DECIDe) threat assesment framework developed by Gary Ackerman and his group (then at the Monterey Institute for International Studies (MIIS)). The DECIDe Framework is based on a "contributing factors approach" that: 1) lays out the key elements (factors) that shape a terrorist group's targeting decision(s); 2) indicates the major relationships and interplay between these factors; and 3) makes clear their direct influences on target selection. The factors and sub-factors used in the framework, as well as the relationships between them, are based upon the conclusions and hypotheses drawn from the literature assessment, case studies of past attacks on critical infrastructure and analysis of data contained in the database of attacks on critical infrastructure. As a first test of our approach, we successfully constructed prototype models of the various MIIS analyses processes and were able to capture the richness of the inter-factor relationships encoded in the DECIDe framework in a computational model that could be used for prediction as well as for training and dissemination.

At the highest level of analysis (see Figure 2), the model consists of a set of classes that model the various factor clusters (such as ideology and psychological factors, group dynamics, past activity, organizational structure and life cycle status, and technical expertise). The clusters are related to each other which are correlations between the different classes. For instance, knowing the ideology of a group can provide some information about the type of target the group is likely to select. Direct correlations are represented as conditional probability distributions which quantify the distribution over values of a child class given an assignment of values to the parent class. The algorithms for inference go backward and forward on the network updating the values of query variables (a subset of the network) given evidence on any another subset of variables.

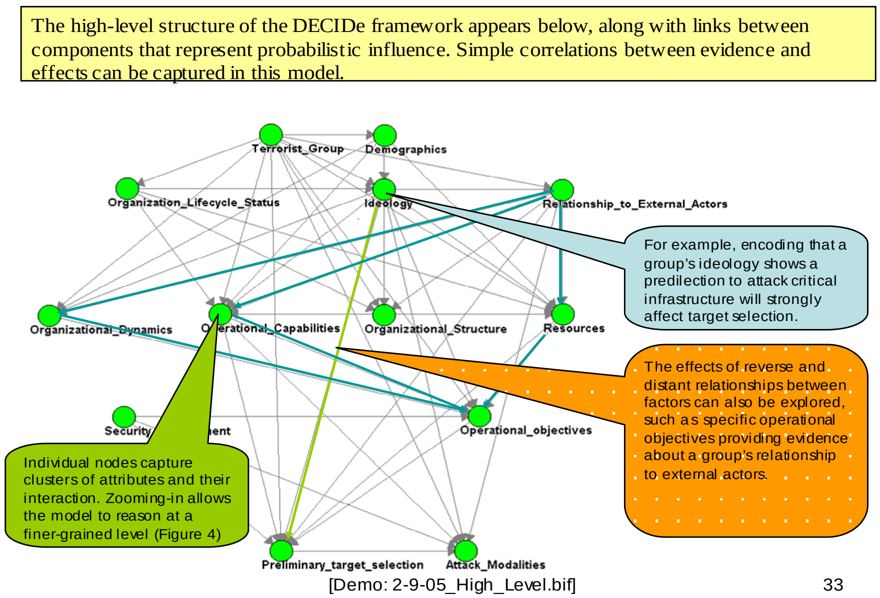

Fig. 2. A formal graphical model of the MIIS threat assessment process. Shown is the highest level of the DECIDe framework with the various factors and the outcome variable (Preliminary_target_selection). Each of these factors represents an aggregation of information from more fine grained analysis (see Figure 3 and Figure 4).

CPRM models are able to reason at multiple levels of evidence and model specificity including detailed process and agent interactions (see Figure 3). The lower level analysis process models the detailed decomposition of individual factors and their inter-relationships. For instance the ideology of a group is composed by their ethnic and political makeup, their attitude towards killing combatants, themselves, or innocent people, etc. The PRM framework allows us to represent and infer with probabilistic relations at multiple levels of granularity. This, we believe, is an essential requirement for dealing with contextual information in any domain. Being able to make use of partially specified or incomplete information requires aggregation of information.

Figure 4 shows the finest level of encoding of the specific events pertaining to an attack taken from the MIIS Critic database. The model is built automatically from the database and provides the analyst a detailed simulation of the events leading to and constituting the attack, and the consequences of the attack. The analyst can query the model in various ways including interactively changing some or all of the events, participants, actions, times, synchronization points, etc and running the modified simulation to make predictions.

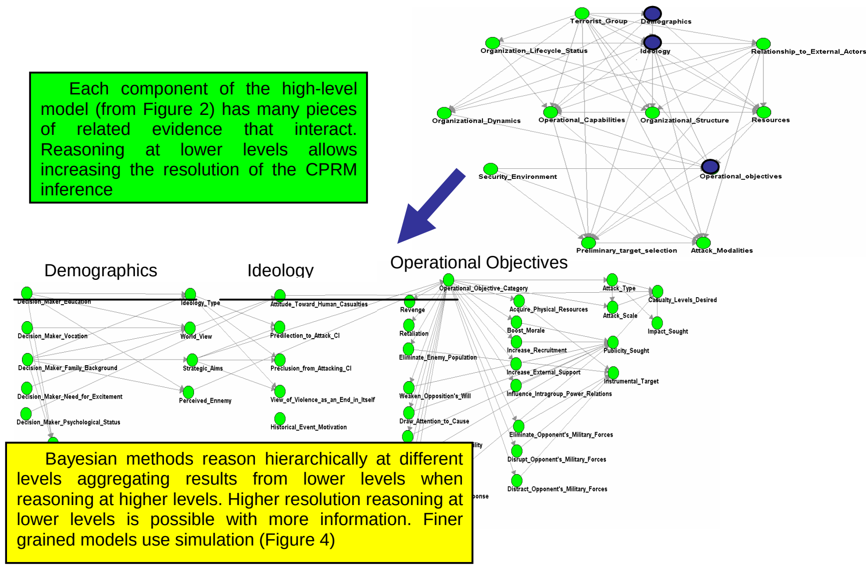

Fig. 3. A lower level probabilistic network for increased resolution of reasoning.

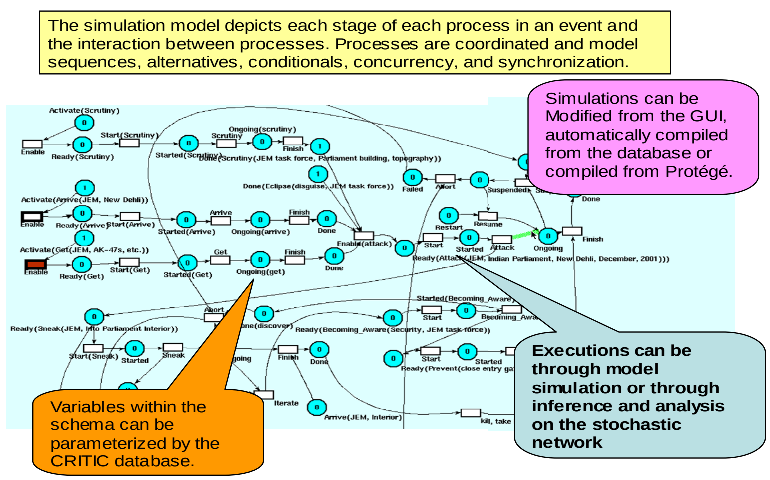

Fig. 4. A fine grained model of event unfoldings taken from parameters of an attack on the Indian Parliament from the MIIS Terrorist Incident DB.

We were able to successfully build the model of threat assessment that faithfully modeled all the factors and their inter-relationships as encoded in the DECIDe protocol developed by MIIS. Besides the obvious advantage of systematizing and making operational the threat assessment process, the model has several features that extend the scope and coverage of the existing model. Important features include the *fine-grained* action and event model that can capture a wide range of possible events and their interactions; b) the *context-sensitive* and *evidential* nature of the model that can model adaptive behavior in a a dynamic and uncertain environment.

It must be emphasized that the results are still preliminary and any model of this level of complexity requires extensive and rigorous evaluation before being fielded. We are planning two specific evaluations. One involves predicting specific cases from the historical database. Our empirically trained, model parameters can be validated by testing it on specific groups and attacks on a gold standard test set (both positive (attack) and negative (non-attacks or thwarted attack) cases). The second, more interesting evaluation would be to compare and cross-validate the model predictions on unknown cases with a trained expert who goes through the manual threat assessment protocol answering specific questions and using the DECIDe aggregation and factor weighting scheme. Any divergence from the empirically derived model predictions could point to incorrect model parameters (maybe due to incomplete data) or conversely to shortcomings and possible improvements in the current DECIDe protocol. Both these evaluations await future results.

4.2 Social and Cultural Factors in Analysis

Social and cultural context is one of the fundamental determinants of human behavior. In general, model building and automated analysis incorporating social and cultural context is at its infancy, partly because systematic delineation and treatment of socio-cultural factors and practices has proven extremely hard. Does the CPRM framework facilitate in any way the encoding and use of social and cultural factors in analysis and prediction?

We took the first tests on testing the utility of our approach by applying the modeling framework to predict the role of cultural and social factors regarding the production and use of coca in Bolivia (especially after election of Bolivian President Evo Morales in December 2005). Specifically, we asked the question of our model "How do cultural, social, and political factors impact coca eradication efforts?"

To answer the question, analyst and anthropologist Katie Sievers came up with a set of relevant *Social/Cultural Factors, Economic Factors, and Political Factors.* We used the framework to model the various factors and to compute the joint impact of the different kinds of factors on coca eradication efforts? We also used the model to predictions differences in responses to coca-eradication efforts between Bolivia and Peru.

	Bolivia	Peru
Education	Adult literacy rate- 87% Primary school attendance- 77.5% Secondary school attendance- 56.5% Source: UNICEF	Adult literacy rate- 88% Primary school attendance- 95.5 % Secondary school attendance- 48% Source: UNICEF
Religion (Catholic, Protestant)	90% Catholic 10% Protestant Source: CIA World Factbook, http://www.providence.edu/las/ Statistics.htm	85% Catholic 10% Protestant other/none- 5% Source: CIA World Factbook,
Ethnicity (Indigenous, mestizo, white)	Quechua- 30% Mestizo-30 % Aymara-25% White- 15% Source: CIA World Factbook	Amerindian Population45% Mestizo- 37% White- 15% Black, Japanese, Chinese & other- 3% Source: CIA World Factbook
Residence (urban/rural)	Urban- 64% Rural- 36% Source: Bolivia National Institute of Statistics (INE)	Urban- 74% Rural- 26% Source: UNICEF
Economic status	Population below $1 per day- 14% Gross National Income per capita- $960	Population below $1 per day- 18% Gross National Income per capita- $2,360 (Note: GNI per capita for the US is $41,400)
Nutrition	Children under 5 suffering from underweight- 8%. Children under 5 suffering from stunted growth (due to poor nutrition)- 27%	underweight- 7% stunted growth (malnutrition)- 25%
Government position re: coca eradication	Current government is strongly against coca eradication. Morales operates on a policy of "Coca yes, Cocaine no"	Garcia government aims to continue coca eradication. Past conflicts between coca growers and the government Coca growers marching on the capital
Political strength of coca growers (cocaleros)	Well organized, fairly strong political constituency.	Well organized, it is unclear how much political clout this group has
Cocalero perception of government attitudes re: coca	Government is working for the cocaleros. Morales a cocalero himself, so having his role in office immediately allows cocaleros to see themselves as part of the govt. However he will have to make good on his promises to reduce eradication in order to keep their support.	Garcia's policy is to continue coca eradications although it is unclear at this point what sort of relationship he will try to create with the coca growers.

Fig. 5. Shown above are some of the relevant factors and activities bearing on Bolivian and Peruvian reponses to a coca eradication policy. Analysts generate a table of this form which is translated into an ontology (using Protege (http://protege.stanford.edu)) which is then automatically compiled into a graphical model for testing, validation, predictions and inference. Notice that some factors have quantitative values while others have qualitative descriptions.

Figure 5 shows a fragment of the information used (culled from literature and on-line sources by anthropologist Katie Sievers) to construct a model for the scenario of interest.

We modeled the different parameters (social, economic, political) by culling evidence from a variety of sources. We instantiated the model, encoded it using the ICSI simulator, and computed the effect of the parameters on the outcome variable (response to a coca eradication policy). Analysts can directly interact with the model GUI and study the effect of perturbations or study the sensitivity of different parameters on the strength of different hypotheses. The modeling framework combines the information and generates the best hypothesis (MAP estimate) for the outcome variables (in this case the support for positive or negative attitudes toward coca eradication programs). Once in the modeling framework, various formal techniques (such as KL divergence, mutual information, sensitivity analysis) can be performed to detect similarities and differences between different social groups/actors/cultures.

Our approach enables us to design and apply formal algorithms for sensitivity and perturbation analysis, divergence in highly sensitive variables including computing ***mutual information and KL-divergence*** between the two structured networks that capture the multiple factors concerning Bolivia and Peru.

Our results predict greater acceptance of coca eradication efforts in Peru compared to Bolivia. While social factors are similar the political factors dominate. Such factors included the different values for the variables {*Leader background, position, attitude}*. An important result of the analysis was the prediction that ***ethnic distribution*** of voters was extremely important as a differentiator. This is initially counter-intuitive since the ethnic distribution of the two countries is fairly similar. However, on closer examination the prediction was made due to the combination of two factors:

1) The *indigenous groups* have a small majority in Bolivia (55%) while they are a large minority (45%) in Peru.
2) *Indigenous groups* are much more likely (odds ratio 4: 1) to vote for coca production than the *mestizo* and *white* groups.

Of course, this is only interesting if we knew the actual voting patterns and percentages in the two country elections. We did not have the information at modeling time and assumed that everyone voted. Whatever the ground truth is in terms of voting percentages, the model prediction is informative since it suggests a potentially important information gathering need. Clearly, if future investigation reveals comparable high voting patterns in the two countries, various ongoing interventions (for example religious and missionary activities or targeted educational policies) could be revaluated and their potential impact on public opinion seen in a new light. One central lesson from this modeling exercise was that it was the combination of factors that is predictive. No single factor is sufficiently predictive.

Our initial results suggest that the technology and graphical modeling approach can encode and evaluate the impact of other socio-cultural factors such as education, religion, ethnicity, nutrition, and economic status. We must emphasize that these results are still preliminary and the problem of systematically comparing the impact of social and cultural context on differential policy responses is a complex and important topic which is still at its infancy. What we found iteresting is that our modeling results are able to indicate that despite the degree of similarity of the above-mentioned factors in Bolivia and Peru, the attitudes toward coca are different. Our results are able to suggest reasons at a deep level for this seeming anomaly. There is ongoing work based on collaborations that is extending these preliminary results to investigate issues of scalability and coverage of the basic approach.

5 Conclusion

Dealing with context is a fundamental requirement in modeling cognitive agents. We are interested in using the modeling framework and simulator to encode a variety of scenarios that incorporate social, economic, and cultural factors, activities, and practice and evaluate the impact of this information on the analytic decision process. In selecting scenarios, identifying relevant socio-cultural processes, and in evaluating the model, we plan to continue our close cooperation with anthropologists and sociologists. Model building, evaluation, and training/dissemination are inherently iterative processes. The CPRM framework enables the analyst to interact with the model in various ways including dynamically (asynchronously at run time) changing some or all of the factors, events, participants, actions, times, synchronization points, etc and running the modified simulation to make predictions. Our ultimate goal is to produce a robust software package and methodology that can enhance the current toolkits available to the modeler (Table 1).

One important effect of having a graphical model is the ability to use parameter estimation techniques to directly estimate the quantitative influence and interdependencies between factors and ultimately the contribution of various factors to the outcome variable (threat potential). This allows a data driven approach to computing factor influence which could be compared to the specific conditional independence assumptions and manually generated weighting schemes in the original model (when there is data).

Evaluation of models continues to be a vexing issue. An obvious measure is the ability of the framework to encode a wide variety of rich scenarios which incorporate systemic, social and motivational context. Other measures include a) incremental knowledge and data reuse in terms of the additional effort required to model and analyze new cases, and b) the ease with which contextual knowledge can be entered, analyzed, and modified using both the ICSI simulator GUI and the OWL-based ontology editors (such as Protégé (using OWL-S)) that directly compile to CPRM models.

Acknowledgments. Thanks to John DeNero for building the first models of the MIIS threat assessment scenarios and protocols. We acknowledge the support of NSF, DARPA under the DAML program, and DTO under the AQUAINT program.

References

1. Anderson, Corin R., Pedro Domingos, and Daniel S. Weld. Relational Markov Models and their Application to Adaptive Web Navigation. ***Proceedings of the Eighth ACM SIGKDD*** *International Conference on Knowledge Discovery and Data Mining* (KDD-2002). 2002.
2. Aziz-Zadeh, Lisa Christian Fiebach, Srini Narayanan, Jerome Feldman, Ellen Dodge, Richard B. Ivry (2006), Modulation of the FFA and PPA by language related to faces and places, *Social Neuroscience* (to appear in 2007).
3. DiMaggio, Paul, Culture and cognition/Annual Review of Sociology/ 23 (1997): pp. 263 – 288.
4. Gedigian, M. Bryant, J., Narayanan, S. and Ciric, B (2006). **Catching Metaphors**, *Scalable Natural Language Understanding Conference, Boston, May 2006.*
5. Feldman, J. and S. Narayanan (2004). Embodied Meaning in a Neural Theory of Language, Brain and Language 89 (2004) pp 385-392, Elsevier Press, 2004**.** http://www.icsi.berkeley.edu/%7Esnarayan/B+L.pdf
6. Fillmore, Charles J. (1968): The case for case. In Bach and Harms (Ed.): Universals in Linguistic Theory. New York: Holt, Rinehart, and Winston, 1-88.
7. Fillmore, Charles J. (1976): Frame semantics and the nature of language. In Annals of the New York Academy of Sciences: Conference on the Origin and Development of Language and Sp 2003 (1980) with Mark Johnson. ***Metaphors We Live By.*** University of Chicago Press. 2003 edition contains an 'Afterword'.
8. Lakoff, George (1987). ***Women, Fire, and Dangerous Things: What Categories Reveal*** *About the Mind* University of Chicago Press. ISBN 0-226-46804-6.
9. Lakoff, George (1989). ***Moral Politics***. University of Chicago Press)
10. Lakoff, George (1999) (with Mark Johnson). ***Philosophy In The Flesh: the Embodied Mind*** *and its Challenge to Western Thought*. Basic Book.
11. Narayanan, S. (1999). Reasoning About Actions in Narrative Undertanding Proceedings of the I*nternational Joint Conference on Artificial Intelligence (IJCAI '99)*, pp. 350-358, Stockholm, Aug. 1-6, 1999, Morgan Kaufmann, San Francisco, CA, 1999.
12. Narayanan, S. (1999). Moving Right Along: A Computational Model of Metaphoric Reasoning about Events Proceedings of the *National Conference on Artificial Intelligence (AAAI '99)*, Orlando, Florida, July 18-22, 1999, pp 121-128, AAAI Press, 1999.
13. Narayanan, S., S. Harabagiu (2004). Question Answering based on Semantic Structures, *International Conference on Computational Linguistics (COLING 2004), Geneva, Switzerland, August 22-29, 2004.*
14. Rosch Eleanor (1983), "Prototype classification and logical classification: The two systems" in Scholnick, E., *New Trends in Cognitive Representation: Challenges to Piaget's Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates: 73-86.
15. Rosch Eleanor (1981) (with C. Mervis), "Categorization of Natural Objects," *Annual Review of Psychology 32*: 89-113.
16. Rosch Eleanor (1978) (with Lloyd, B., eds). *Cognition and Categorization*. Hillsdale NJ: Lawrence Erlbaum Associates
17. Pfeffer, Avi (2000) Probabilistic Reasoning for Complex Systems, A.J. Pfeffer. PhD Thesis, Stanford University, January 2000
18. Scheffczyk, Jan, Collin F. Baker, Srini Narayanan (2006), Ontology-based Reasoning about Lexical Resources, *Ontology and Lexical Resources in Natural Language Processing, Cambridge University Press, 2006 (to appear)***, earlier version presented at *OntoLex 2006***, *Genoa, Italy, May 2006.*