Estimating the Shape of Covert Networks

Matthew Dombroski
Paul Fischbeck
Kathleen M. Carley
Carnegie Mellon University

Contact:
Prof. Kathleen M. Carley
Institute for Software Research International
Carnegie Mellon University
Pittsburgh, PA 15213

Tel: 1-412-268-6016
Fax: 1-412-268-1744
Email: kathleen.carley@cmu.edu

Modeling and Simulation
  (Dombroski - student)

Estimating the Shape of Covert Networks

Abstract:
Social network analysis has been used to understand groups of individuals and how they operate. Most of the literature in social networks has dealt with overt organizations with an easily discernable network structure. This paper examines the possibilities of using the inherent structures observed in social networks to make predictions of networks using limited and missing information. The model is based on empirical network data exhibiting the structural properties of triad closure and adjacency. Triad closure indicates that if person i has a dyad with person j and person j has a dyad with person k, then there is a higher than chance likelihood that person i and person k have a dyad. Adjacency is a corollary of triad closure stating that if person i has a dyad with person j, it is more than likely that person i has a dyad with person k. The model exploits these properties using an inference model to update adjacent dyads given information on a reference dyad. The model is tested against several networks to understand and discern its behavior. The paper illustrates that if the model is built with careful consideration towards the network being predicted, it may assist in making better decisions regarding uncertain organizational phenomenon. However, the model performs relatively poorly if there is a disproportionate amount of information either supporting or not supporting a dyad and/or if dyadic priors are well informed. The method is applied in a covert network example, and has been extended for epidemiological networks and improving performance in organizations operating under stress. The paper opens up new avenues in the development of models designed to make network predictions and use those predictions to make better decisions.

Unknown Network Structures

TRADOC PAM 525-5 (1994) hypothesizes that Command and Control will face new and unconventional threats in the post Cold War environment. These threats, largely a result of increased global instability and the rise of regional conflicts, require a change in the tactics and techniques of Command and Control. Arquilla and Ronfeldt (2002) indicate that future conflicts will take place against asymmetric threats consisting of networked forms of organization. These networked forms of organization would have the ability to cloak their activities from detection using dispersed organizational forms and swarming tactics. The events of September 11, 2001, the USS Cole bombing, and embassy bombings in Africa by terrorist cellular units illustrate the effectiveness of networked forms of organization and swarming tactics in the post Cold War era. Such threats require sophisticated tools and technologies to coordinate information and construct an accurate picture of threats and the particular risks they pose for the United States and its Allies. Such techniques must recognize the multifarious nature of such threats, both in the organizational structure of such terrorist threats, and the particular skills, tasks, and resources characteristic of individuals in these groups (Krackhardt and Carley, 1998).

Social network analysis has been used to understand organizational dynamics in a variety of application areas (e.g., epidemiology, technological diffusion, and management consulting). A group's behavior, values, and/or performance can be articulated by understanding the relationships that exist within the group. Most applications to date have been on open groups or societies in controlled experiments. Currently there have been very few network applications to covert or "hidden" networks of interest. Sparrow (1991) discusses the prospects for using social network analysis as investigative tools for intelligence and law enforcement. Sparrow discusses the prospects of using traditional social network centralization measures to identify key individuals in a covert organization and inferring their activities through their connections. Erickson (1981) and Baker and Faulkner (1993) discuss the structure of covert organizations and their distinguishing structural characteristics. Baker and Faulkner conclude that the requirement for secrecy distinguishes the covert organization from the overt organization, which permeates every aspect of the organization including its structure and productivity. Covert organizations with high task loads that require coordination are generally more hierarchical than organizations with lower task loads and less coordination, although all covert organizations studied exhibited flatter and more dispersed forms of organization than comparatively sized overt organizations. PCI (2002) has developed software and tools to assist law enforcement professionals to coordinate information gathering techniques on illicit organizations and their structure through the use of Anacapa charts and link analysis to build social networks. The FBI's Big Floyd is a template matching system designed to link archived criminal tactics and organized criminal groups with current criminal investigations (Bayse and Morris, 1987). Although the tools and techniques described up to this point would be useful to Command and Control for detecting and understanding the activities of networked criminal and terrorist organizations, a tool that can infer the structure of these covert organizations using known social network properties and limited observational data does not currently exist. Such a tool would not only be beneficial in building a more complete picture of covert networks using limited data, but could also allow policy makers to make better-informed decisions.

This paper presents an empirically based probabilistic model, grounded on observational social networks, to infer network structure using limited and incomplete information. First, relative similarity information is used to build a prior probability assessment of who communicates with whom. As direct information on dyadic likelihood is received, these priors are updated. Adjacent dyads are updated through an empirically based inference model. This continues until the likelihood of every dyad in the probabilistic network is inferred and updated. The resulting network provides an estimate of the actual network and may be used to guide policy analysis.

Network Properties

Researchers have uncovered inherent structural properties in social networks (Skvoretz, 1990). These properties arise from the structure of the network itself and not from the behavior of the individuals in the network. They include reciprocity, triad-closure, and triad-closure reciprocity. A corollary of the triad properties is an adjacency property. Simply stated, if persons i and j are talkative with each other, then they are likely to be talkative with others. Formally, if A and B are adjacent dyads, then

if $n_B > \bar{n} \Rightarrow E(n_A) > \bar{n}$, and

if $n_B < \bar{n} \Rightarrow E(n_A) < \bar{n}$,

where n_B is the number of interactions recorded on dyad B, $E(n_A)$ is the expected number of interactions on dyad A, and n-bar is the mean number of interactions for the whole network. In other words, if B has above average activity then the expected value of the distribution of interactions for all of its adjacent dyads will also exceed the mean number of interactions. The degree to which these properties exist varies from network to network (Krackhardt, 1987).

Constructing the Model

The problem domain will determine the relationship of interest (ROI). In most real-world situations, only samples of interactions between individuals can be observed. Depending on the type of interaction, knowing that i and j interacted will inform our belief about the likelihood of a ROI existing between the individuals. But, what, if any, inference can be made about these individuals' relationships with others in the network?

For illustrative purposes and to facilitate model development, we focus on one social network dataset, Bernard and Killworth's 1979 observed interactions between 58 fraternity brothers at a West Virginia university. Because of the size of this data set, it was not possible to develop a robust inference model based on the triad-closure property. Instead, the model is based on adjacency properties found in the data. Figure 1 shows this relationship between interactions on a reference dyad and the expected number of interactions on an adjacent dyad. As the number of communications for the reference dyad increases, so does the expected number for the adjacent dyads.


Figure 1—Adjacency Property Illustrated using Bernard and Killworth Fraternity Data

Transforming Data

In order to build a model of dyadic dependency, the network data of "counts of interactions" has to be converted into probabilities that an ROI exists between any pair of individuals. To do so, requires a careful definition of what constitutes an ROI. Two important considerations must be made.
1. The number of interactions needed to define when a ROI exists
2. The marginal increase of each additional interaction towards the probability of a ROI existing

We also need to establish the functional form that relates additional interactions to the probability of ROI. In this paper, we are assuming a concave function (marginal decreasing value). A standard exponential functional form is used:

$$d_{ij} = \frac{(1 - e^{-\lambda x_{ij}})}{(1 - e^{-\lambda (\max(x_{ij}))})}$$

where d_{ij} is the probability of the relationship of interest, λ is the shape parameter of the function (higher values more concave), x_{ij} is the interaction dyadic data, and max (x_{ij}) is the ROI threshold value. For this model, a ROI threshold of 21 interactions and a λ value of 0.14 were used. The 21-interaction threshold value was chosen so that strong relationships could be modeled and that a large distribution of reference probabilities could be considered. The 0.14 λ value was chosen as a moderate value to attain some concavity to the curve.

Building a Model of Dyadic Dependency

The dependency relationship between dyads can be illustrated by plotting adjacent dyads' probabilities against reference dyads probabilities, for all dyadic pairs. Figure 2 shows the percentile contours for the transformed fraternity data. Note that the percentile contours are generally increasing with probability.
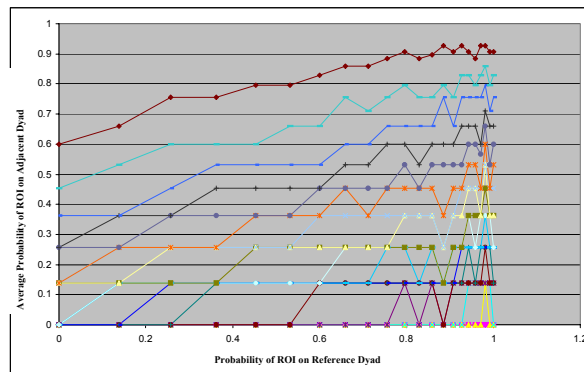

Figure 2---Raw Fraternity Data Showing the Dependency Relationship Between Dyads

Fitting lines through these data and smoothing the parameters can clearly show the dependency between dyadic probabilities in the fraternity data. For this application, the relationships shown in the plots were modeled using a neural network. (Freeman and Skapura, 1991) The neural network consists of 9 nodes and 3 layers.

Implementation

Priors to dyads may be assigned using homophily (McPherson and Smith-Lovin, 1987, McPherson, Popielarz, and Drobnic, 1992, Valente et al., 1997) and a database of social relationships and attributes such as the PCANS methodology (Krackhardt and Carley, 1998). As an observation comes in to inform the model, Bayes Rule performs the direct update and the inference model can then propagate the information to update other relationships in the network. Suppose that I_{ij} is the event that an interaction is observed between nodes i and j. Suppose also that L_{ij} is the event that the ROI exists between nodes i and j. P $(I_{ij}|L_{ij})$ and P $(I_{ij}|L_{ij}{}^C)$ can be assessed for each piece of incoming information. Then the conditionals can be used to update the probability of the reference dyad, P (L_{ij}).

$$P(L_{ij} \mid I_{ij}) = \frac{P(I_{ij}L_{ij})}{P(I_{ij})} = \frac{P(I_{ij} \mid L_{ij})P(L_{ij})}{P(I_{ij} \mid L_{ij})P(L_{ij}) + P(I_{ij} \mid L_{ij}{}^C)P(L_{ij}{}^C)}$$

Once the probability of the initial dyad is calculated using Bayes Rule, there are several choices for how to propagate the update through the rest of the network. In this paper three alternative models are considered:

1. Bayes Rule is used to update only the reference dyad.

2.   Bayes Rule is used to update the reference dyad. A secondary round of updates are applied to the dyads immediately adjacent to the reference dyad, using the inference model.

3.   Bayes Rule is used to update the reference dyad. A secondary and tertiary round of updates are applied to the dyads immediately adjacent to the reference dyad and adjacent dyads using the inference model.

Simulation and Analysis

In order to test the model implementations and understand its effectiveness under different parameters, the models were simulated on a sample 20 node network from the fraternity data set. The following parameters were varied in the simulations.

- Input Data Accuracy
- Prior Assessments
- Proportion of Information Supporting Existence of Dyads

The metric chosen to evaluate the simulations is absolute error, which is defined below:

$$\text{Absolute Error} = (\sum_{i>j}^{N} \sum_{j=1} abs(p_{ij} - A_{ij}))$$

where p_{ij} is the predicted probability for the dyad between nodes i and j and A_{ij} is the actual probability for the dyad between nodes i and j. Simulations consisted of 20 runs under each condition where an interaction was sampled from the real network of interactions and used to build a prediction. Positive and negative dyadic data were sampled; where it is assumed that using 1-Aij can generate negative dyad data. Each simulation consisted of 1000 updates, although the graphs below only plot the first 250 updates. Significance of differences between the models was calculated using a paired t-test at the 0.05 level. A base case was assumed with the following parameters from which deviations were made in the input parameters in later simulations:

- 0.2 uninformed uniform probability for each dyad
- 0.5 probability of receiving updates supporting a dyad
- P ($I_{12}|L_{12}$) = 0.56, P ($I_{12}|L_{12}^{C}$) = 0.24

Using an informed prior, model 3 performed significantly worse than model 1 across all updates and models 1 and 2 were indistinguishable. This is not a surprising result since the network is essentially already updated when an informed prior is used. Since models 2 and 3 affect all dyads with each update, erroneous priors get corrected much more quickly than for model 1. Model 3 would be expected to converge quicker to the real network since it changes more priors than model 2. However, the result indicates that model 3 might be over-inferencing the network, making it less effective than model 2.
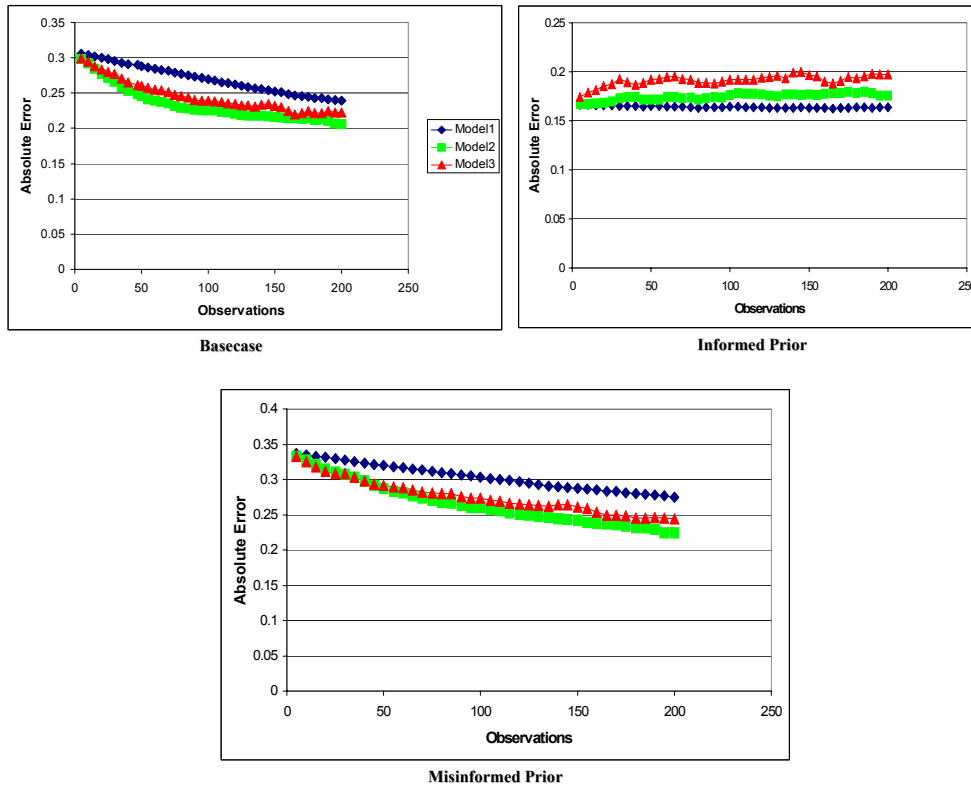
Figure 3---Simulation results by varying the prior assessments for different network prediction models showing that inference outperforms the control in all cases except an informed prior

It was found that the models' results are sensitive to the proportion of updates supporting the existence of the ROI for a dyad. In both extreme cases where the probability of receiving a supporting update was 0.2 and 0.8, models 2 and 3 performed well in the first 50 updates, but were quickly surpassed by model 1 thereafter. Because an imbalanced number of supporting (non-supporting) updates are arriving, the models that use the adjacency property amplify this imbalance and drift to over (under) predict. It appears that these models work well when there is large network uncertainty and limited data. At some point it would be optimal to switch from the inference models to model 1, but knowing when to do this would likely depend on several unknowable properties of the network. Future research will examine this scenario in greater detail.

Basecase

0.8 P(observing dyad supporting information)
0.2 P(observing non dyad supporting information)

0.2 P(observing dyad supporting information)
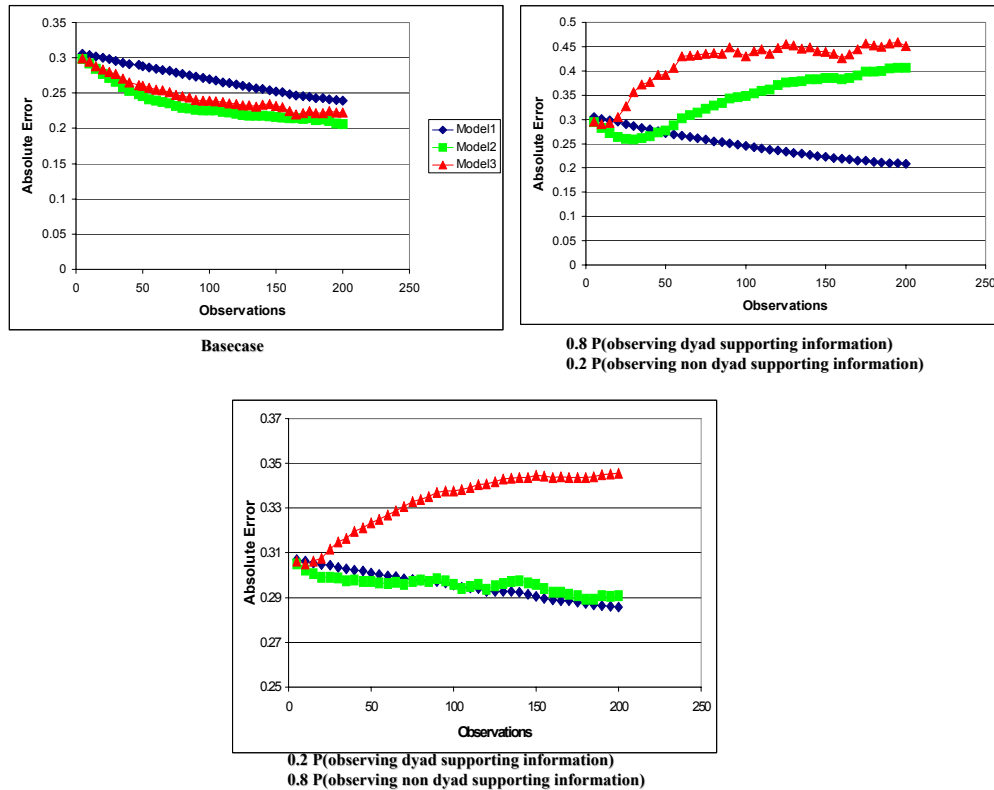**0.8 P(observing non dyad supporting information)**

Figure 4---Simulation results by varying the proportion of incoming data that supports the existence of a dyad illustrating that the inference models over predict the network when a disproportionate amount of information comes in

Changing the conditionals to 0.8 and 0.1 resulted in model 2 outperforming model 1 in the first 50 updates, but the results are not significant. With random conditionals the inference models perform significantly better than model 1 for the first 250 updates. This is an important result and leads to the conclusion that the inference could be effective when operating under uncertain conditions. Highly accurate data will not distinguish the performance of the models.
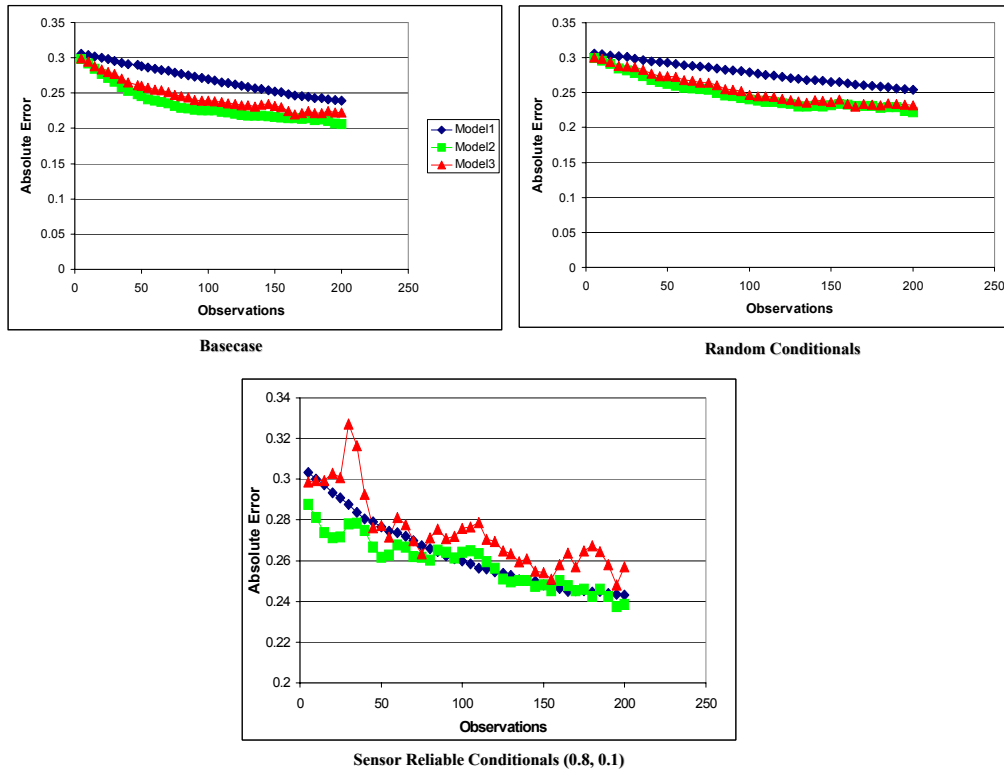
Sensor Reliable Conditionals (0.8, 0.1)

Figure 5---Simulation results by varying the reliability of the data by varying conditionals illustrating that the inference models perform well when conditionals are random, but do not perform well when very reliable

Decision Analysis-Covert Networks

Consider a subpopulation of conspirators is being examined to determine its network structure and who to target in that population with an unknown network structure. Borgatti (2002) identifies a key player metric for network analysis that we use to develop inoculation strategies. Suppose that resources are available to remove 7 out of the 20 individuals. We can examine the efficacy of the recommended strategy by comparing the recommended strategy (developed from a network prediction after 100 updates) to that of the real network. Minimums, averages, and maximums for the 20 simulation runs under models 1, 2, and 3 are plotted in Figure 3.


Figure 6---A comparison of models for covert networks; for all possible isolations, model 2 has a higher average value for the number of fragments created

Table 1-- Evaluation of models for covert network scenario showing the expected values of fragments, given different models, and EVPI

		Fragments	
	total	isolates	subgroups
perfect information	12	11	1
E(Model1)	4.9	3.45	1.45
E(Model2)	5.4	3.6	1.8
E(Model3)	4.9	3.55	1.35

Figure 5 and Table 1 indicate that on average, model 2 is best at breaking up the network, but model 3 attains a maximum number of fragments that is higher than model 2 for isolates and total number of fragments. Comparing these values to what they would be with perfect information we see that model 2 performs best of all the models, but none of the models perform as well as perfect information. None of the models for any runs attain the perfect information values. The results in Table 4 indicate that on average, model 2 will result in 0.5 more total fragments than either model 1 or 3. But there is uncertainty, as choosing model 3 might yield a better result than model 2.

Conclusions

The simulation results show that the inference models' performance is mixed. The simulations were run against a portion of the data that the model was built from, meaning that the results shown are a best-case scenario. The models do perform relatively well when there is a proportionate amount of information supporting and not supporting the existence of a dyad. If there is a disproportionate amount of information in either direction, the model has a tendency to over or under predict the network. Future work will examine whether or not misinformed network predictions from over or under informed networks leads to poorer decisions by applying those network predictions to the decision component. The models also seem to perform well when priors are either uninformed or possibly misinformed. Future work will explore why the inference models perform relatively poorly when priors are well informed. It is likely that with informed priors, one may not need to use inference to inform the network and the effect shown in this paper is the result of well-informed priors being manipulated by inference to become less informed. The decision analysis illustrates that the inference models may lead to better decision-making than the control model, although the improvement in decision-making might be relatively small. It must be emphasized that this is a best-case scenario and as a result if this model is to be applied on uncertain networks, then a great deal of care must be taken to ensure that the network used to build the model must be very close to the uncertain network that is being predicted. In addition, the model tends to over infer dyads after 300 updates in all simulation scenarios and as a result, care must be taken to consider when to stop using the inference model to prevent systematic errors in prediction. Nevertheless, keeping these principles in mind, paying close attention to the network used to build the model, and paying close attention to the distribution of input information used to inform the network and infer or not infer dyads, the models have the potential to improve decision-making under uncertainty.

Future work will examine the model in greater depth and integrate new prediction metrics and tools into the model, such as new biases, both structural and behavioral, to increase the effectiveness of the prediction. Homophily, or the tendency for individuals exhibiting similar

characteristics to be in contact with each other, provides a useful behavioral bias that could be integrated into the model in the future to improve network predictions under uncertainty. The model indicates that using the adjacency property for network prediction likely leads to a modest improvement in network prediction versus the alternative of not inferring structure. However, a more powerful property, such as triad-closure, would likely lead to better prediction results. Unfortunately, current datasets do not provide the density required for a rigorous statistical analysis needed to construct a similar triad-closure model.

An important conclusion to draw from this analysis is that different organizations exhibit different network properties to different degrees. Because the results from this analysis are mixed, a great deal of care must be exhibited when choosing a sample dataset with which to build this model from and use it for network prediction. Specifically, the social network literature identifies work related organizations as exhibiting different properties from social related organizations. Terrorist and organized crime organizations likely exist as a hybrid between work and social oriented organizations in their structure, while also exhibiting their own unique characteristics (Krebs, 2001). We can hypothesize of a suite of models, each one modeling a different organization or set of organizations that may be used for network prediction, such as work-team relationships, friendship relationships, and project coordination relationships. Each model could be used to characterize different organizations or the same organization, but from different perspectives, which could offer insight into the organization and its functions, activities, and resources. We can also use such models to understand the evolution of an organization and the collective understanding of that organization. As more and more information about a covert organization becomes available, different models can be employed to predict the structure of the organization as Command and Control becomes more familiar with the organization and its structure.

This study has also indicated that although gaining insight about the organizational structure of a covert organization is important, it should not supersede the requirement for traditional information gathering techniques that focus on the resources and tasks that characterize individuals in the organization. Unfortunately, network construction techniques and tools such as this one, may put a strain on valuable investigative resources that could be used for other investigative tasks. Future work will link network prediction tools, such as this, with multi-agent technology to simulate covert organizations, their information diffusion behaviors, and illicit activities. This research will focus on destabilization techniques to determine whom to isolate in a covert organization under uncertainty. If the objective is to destabilize and break an organization apart, removing the most central individual (from a social network perspective) in the organization may not be the most effective strategy.

Future work will continue to explore the decision scenarios that this model can be applied to. As shown in the paper, the simulation results translate into better overall decisions by breaking the network apart into more pieces by inferring dyads than not inferring dyads. Future work will expand on the decision analysis component, not only applying the model to other decision contexts, but exploring the space where improved decisions can be made and where poorer decisions can be made. Specifically, an analysis of the effectiveness of destabilization will be examined at different numbers of updates to determine the point at which in a decision context a decision maker becomes indifferent between inferring dyads and not inferring dyads. This study will compare the decision analysis results back to the simulation results to see if there is a correlation between absolute error in the simulation and the effectiveness of the predicted

destabilization strategies. Such an analysis will improve assessment techniques of simulations in the future, relating the results back to the required decision context.

References

Arquilla J. and D. Ronfeldt (Eds.), 2001, *Networks and Netwars: The Future of Terror, Crime, and Militancy*, RAND Publication, National Security Division, available online: http://www.rand.org.

Baker, W.E. and R.R. Faulkner, December, 1993, "The Social Organization of Conspiracy: Illegal Networks in the Heavy Electrical Equipment Industry." *American Sociological Review* **58,** pp. 837-860.

Bayse, W.A. and C.G. Morris, 1987, "FBI Automation Strategy: Development of AI Applications for National Investigative Programs," *Signal Magazine*.

Borgatti, S.P., 2002, "The Key Player Problem," Proceedings from National Academy of Sciences Workshop on Terrorism, Washington DC.

Erickson, B, September, 1981, "Secret Societies and Social Structure," *Social Forces* **60**(1), pp. 188-210.

Freeman, J.A. and D.M. Skapura, 1991, Neural Networks: Algorithms, Applications, and Programming Techniques, Addison-Wesley Publishing Company, Mass.

Killworth, P.D. and H.R. Bernard, 1979, "Informant Accuracy in Social Network Data III: A Comparison of Triadic Structure in Behavioral and Cognitive Data," *Social Networks*, **2**, 10-46.

Krackhardt, D., 1987, "Cognitive Social Structures," *Social Networks*, **9**, 104-134.

Krackhardt, D. and K. Carley, 1998, "A PCANS Model of Structure in Organization," In: *Proceedings of the 1998 International Symposium on Command and Control Research and Technology*, 113-119.

Krebs, V., 2001, "Mapping Networks of Terrorist Cells," *Connections*, **24**(3), 43-52.

McPherson, J.M., L., Smith-Lovin, 1987, "Homophily in Voluntary Organizations: Status Distance and the Composition of Face-to-Face Groups," *American Sociological Review*, **52**, 370-379.

McPherson, J.M., P.A. Popielarz, and S. Drobnic, 1992, "Social Networks and Organizational Dynamics," *American Sociological Review*, **57**, 153-170.

Precision Computing Intelligence (PCI), 2002, "Crimelink: Investigative Analysis Software," Precision Computing Intelligence, available online: http://www.crimelink.com/, Sierra Vista, AZ.

Skvoretz, J., 1990, "Biased Net Theory: Approximations, Simulations, and Observations," *Social Networks*, **12**, 217-238.

Sparrow, M.K., 1991, "The Application of Network Analysis to Criminal Intelligence: An Assessment of the Prospects," *Social Networks*, **13**, 251-274.

TRADOC Pam 525-5, August 1, 1994, "Force XXI Operations," Department of the Army, Headquarters, United States Army, Training and Doctrine Command, Fort Monroe, Virginia.

Valente, T.W., S.C. Watkins, M.N. Jato, A.V. Straten, L.M. Tsitsol, 1997, "Social Network Associations with Contraceptive Use Among Cameroonian Women in Voluntary Associations," *Social Science Medicine,* **45**(5), 677-587.

Appendix-Significance Plots of Differences Between Absolute Errors of Model Simulation Results

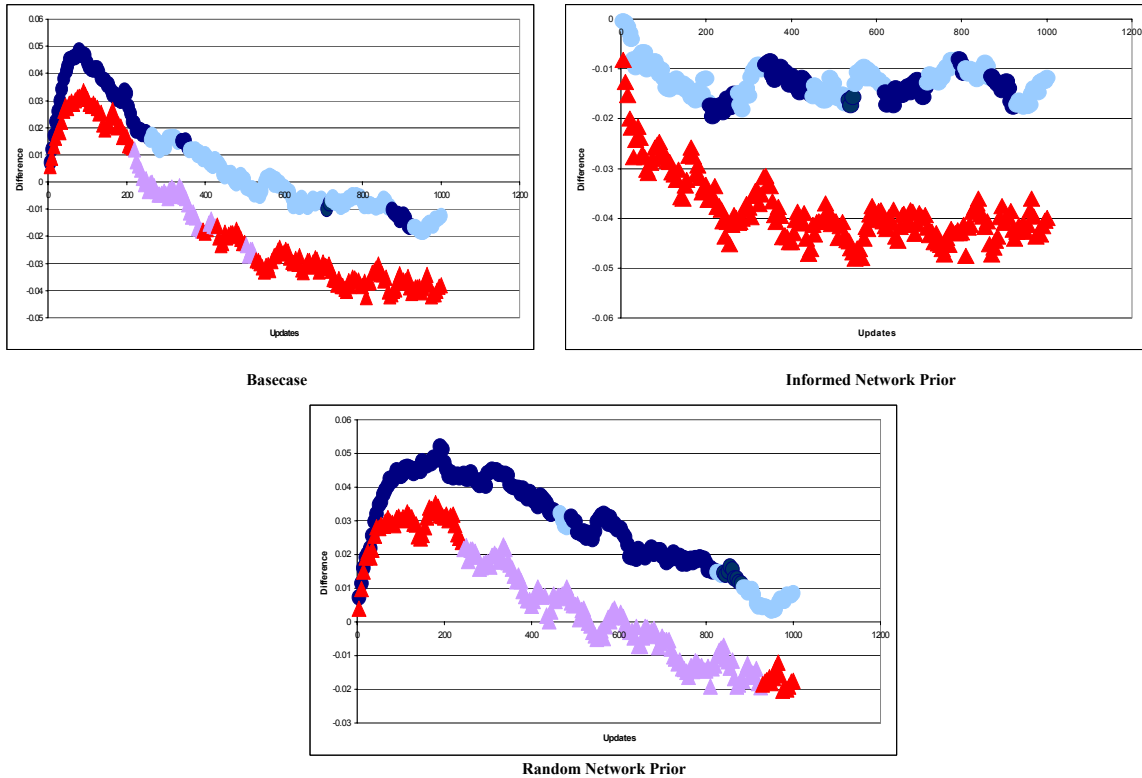

Basecase

Informed Network Prior

Random Network Prior

Figure 7---Absolute Error Statistical Comparison Plots for Different Network Priors— Bold circles-Significant (0.05 level) differences of model 1 (no inferencing) minus model 2 (inferencing on immediately adjacent dyads) absolute error---Standard circles-Non-significant differences of model 1 minus model 3 (inferencing on entire network) absolute error---Bold triangles-Significant (0.05 level) differences of model 1 minus model 3 absolute error---Standard triangles-Non-significant differences of model 1 minus model 3
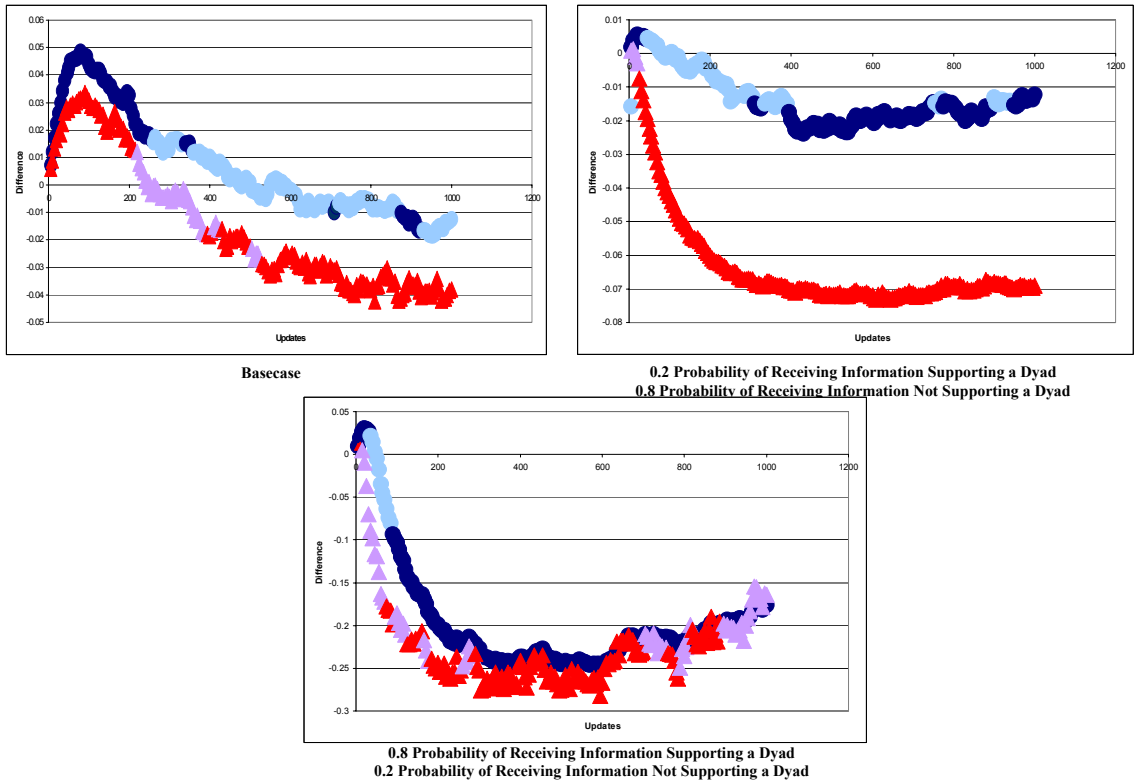
Basecase

0.2 Probability of Receiving Information Supporting a Dyad
0.8 Probability of Receiving Information Not Supporting a Dyad

0.8 Probability of Receiving Information Supporting a Dyad
**0.2 Probability of Receiving Information Not Supporting a Dyad**

Figure 8---Absolute Error Statistical Comparison Plots for Different Probabilities of Supporting Information— Bold circles-Significant (0.05 level) differences of model 1 minus model 2 absolute error---Standard circles-Non-significant differences of model 1 minus model 2 absolute error---Bold triangles-Significant (0.05 level) differences of model 1 minus model 3 absolute error---Standard triangles-Non-significant differences of model 1 minus model 3

Basecase

Randomly Assigned Conditional Values
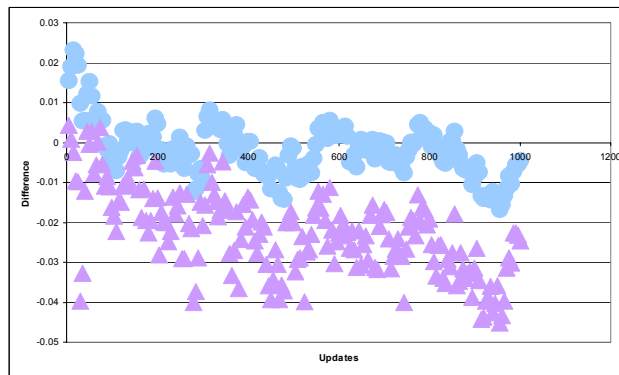
$0.8\ P(I_{ij}|L_{ij})$ and $0.1\ P(I_{ij}|L_{ij}{}^C)$

Figure 9---Absolute Error Statistical Comparison Plots for Different Conditional Probabilities— Bold circles-Significant (0.05 level) differences of model 1 minus model 2 absolute error---Standard circles-Non-significant differences of model 1 minus model 2 absolute error---Bold triangles-Significant (0.05 level) differences of model 1 minus model 3 absolute error---Standard triangles-Non-significant differences of model 1 minus model 3