13th ICCRTS: C2 for Complex Endeavors

Automated Influence Network Generation and the Node Parameter Sensitivity Analysis

C2 Assessment Tools and Metrics

Il-Chul Moon, "STUDENT", Eunice J. Kim, and Kathleen M. Carley

Point of Contact: Il-Chul Moon

Carnegie Mellon University

5000 Forbes Ave. Wean Hall 1325, Pittsburgh, PA, 15213

412-268-5866

icmoon@cmu.edu

eunicek@andrew.cmu.edu

carley@cs.cmu.edu

Automated Influence Network Generation and the Node Parameter Sensitivity Analysis

Abstract

An influence network is a directed graph extensively used for Effects-Based Operation. It contains nodes that represent events and links that encode causal relationships among events. It propagates the likelihood of each event through promotion or inhibition by its parents. As a subject matter expert often builds this network by hand, we helped simplify the influence network generation in Organization Risk Analyzer. The resulting influence network is generated from a multi-mode, multi-plex organizational network structure, and the generation scheme is based on assessing event flows and evaluating the factors on task management of the organization. To support the soundness of such network generation, we provide sensitivity analysis of baseline probabilities, a major parameter of the model, by bootstrap sampling of the leaf nodes and propagating different levels of assigned parameters. Finally, we provide an example of analysis by utilizing the introduced generation method and a dataset from 1998 US embassy bombing in Kenya.

Keywords: Influence network, Effects-based operation, Sensitivity analysis

1. Introduction

An influence network (Wagenhals and Levis, 2007) is a directed graph used to estimate the likelihood of events. It contains nodes that represent events and links that encode causal relationships among events. It propagates the likelihood of each event through promotion or inhibition by its parents. In the real world, the influence network is becoming popular, as knowing how to influence and redirect the change of situation is becoming important. For instance, Wagenhals and Levis (2007) designed an influence network focused on subduing IED attacks in Iraq. Hudson et al (2001) introduces potential usages for counterterrorism, and Rosen and Smith (1996) show an influence network model for building a military and diplomatic strategy. The influence network contains belief statements related to politic, military, social, economy, information and infrastructure, so called PMESII (DARPA, 2005; Silverman, 2007) in military planning. The network helps evaluating on which sector friendly forces should act to lower the IED attack frequencies in the region (Hufbauer et al, 2001). This approach is different from the traditional action-based operation, which focuses on sweeping regions, setting up multiple checkpoints, and ignoring the cultural and sociological consequence of such actions.

This paper introduces a framework generating such an influence network evaluating the task completion likelihood of a key task. In the context of our problem, we have belief statements about personnel sufficiency, resource availability, information accessibility, organizational support, and etc that influence the completion of a certain task. In addition to a single task analysis, these sub tasks required by a final task are interwoven to others in a task network. Therefore, the result of a final major task is influenced by a set of subtasks. Thus, the prior tasks in the task network and accompanying belief statements for each of these tasks create a big picture of influence network resulting in the final event occurrence likelihood. On the other hand, we already have an organization structure in a meta-network format where we can infer the above task completion factors as well as a task network. Thus, we create a function that generates an influence network from the social network in completing the given task.

In this paper, we use the above influence network generation idea and implement a function that automatically generates an influence network regarding a terrorism act, e.g. the 1998 US Embassy bombing in Kenya. This automatic generation assesses the likelihood of adversarial operations' success, and the generated influence network contributes to determining friendly forces' optimal course of action that reduces the enemy's success likelihood.

2. Previous research

Previous influence network generation is done by human subject-matter experts. These traditional creation approaches make the influence network subject to the experts' prejudice, specialty, and so on. Furthermore, the experts need significant time to create an influence network. In this section, we discuss such short-comings of the traditional generation. Also, we compare the influence network generated by our approach to the ones produced by the traditional approach.

2.1. Traditional and automatic generations of an influence network

Traditionally, influence network has been produced by hands of subject matter experts. They have knowledge of the target situation and organization, assess belief statements related to a target event or effect, and draw an influence network by setting up its nodes, links and parameters based on their own knowledge. However, this creates a number of problems in real usage of this inference tool. First, the generation takes a long time. Second, the generation is sometimes too subject to experts' opinions (Vego, 2006) . Currently, experts decide on what related belief statements are, how the topology shapes the linkage of beliefs, what the baseline probability of the each belief should be, and etc. However, without a template or a commonly accepted practice of network generation, the influence network created would be easily biased by an individual analyst. Therefore, we need a tool that creates a blueprint of an influence network with a standardized template that experts can examine and customize based on their expertise.

To automatically generate an influence network using a template, we need to resolve two issues: the representation of experts' knowledge about target organizational structure; and which viewpoints the tool should take to assess the organization with the structure. Some analysts suggest that an organization's task performance level is based on its structure. For instance, loose network (Burke, 2004; Hoffman, 1998) is a term describing a decentralized yet effective structure. Thus, we use meta-network concept (Carley, 2006; Krackhardt and Carley, 1998) to examine the organizational structure critical to operational environment. A meta-network is an extended version of a social network including various elements of an organization. For instance, it represents a task assignment by linking an individual to a task. In this manner, resource and information distribution, work relations and task dependencies are all captured in a meta-network representation. Thus, we are going to use a meta-network as means to represent a current organizational situation. This meta-network is a product of data-mining from open-source information, expert knowledge on a target organization, and existing relational intelligence. Therefore, it will be able to cover broader range of domain knowledge compared to an expert's knowledge only.

While using a meta-network as a domain knowledge representation format, the generation function should examine factors affecting the situation or event completion likelihood. Since we limit the scope of its application to the estimation of task completion likelihood, the function looks at the factors affecting whether a task is completed or not. From the organization management perspective, a list of factors for task completion has been enumerated. This list is not complete, but we believe that it captures most of the salient features regarding the completion of a task. The list includes *prior task completion*, *personnel assignment*, *task importance*, *task complexity* and *resource/expertise availability*. Our list emerged from the organizational management and operations research domains. Researchers have identified factors derived from the nature of tasks or organizational structure. For example, the operations research domain has developed task precedence network analysis (Eisner, 1962). It suggests better ways to organize the task performance plan or to minimize the impact of completion delays, etc. While the task dependency is one factor considering the links among the tasks, the task complexity and the importance of each task are other factors that affect tasks completion (Campbell, 1991; Forsyth and Schlenker, 1977). Organizational structure suggests the criticality of personnel, resources and information distribution. Human resource management is another approach to enhance the organizational performance by assigning personnel to tasks effectively (Becker and Gerhart, 1996). Furthermore, as organizations perform knowledge-intensive tasks, the diffusion of knowledge or knowledge management becomes another important factor in getting a job done (Argote and Ingram, 2000).

2.2. Differences among meta-network, traditionally created influence network, and automatically generated influence network

To illustrate an automatically generated influence network in detail, we compare it to meta-network and traditionally created influence network. Through the comparisons, we show what should be inferred and assumed to fill the gap between the meta-network and the influence network. We organized the comparisons in Table 1.

As in Table 1, the meta-network and the influence network have different meanings in their nodes and links. For instance, the nodes in a meta-network are entities while those in an influence network are belief statements. Therefore, we combine a set of linked entities and infer the relations among them, and we produce belief statements out of these relations. This approach is similar to the narrative network representation (Pentland and Feldman, 2007) that stores a story of operations in a network formation.

Additionally, nodes and links in a meta-network do not have parameters except edge weights showing the strength of the link. However, an influence network requires three parameters: baseline probabilities for nodes; and inhibition and promotion parameters for links. Whereas traditionally created influence networks obtain these values from subject matter experts, we supply these values by utilizing a set of heuristics assessing the situation and assigning predefined baseline probabilities. As far as concerning the inhibition and promotion parameters, we use default values.

Table 1: The comparison of meta-network, automatically generated influence network, and traditional influence network

	Meta-Network	Automatically generated influence network	Traditionally created influence network
Node	Entities in an organization	Belief statements in a predefined template	Belief statements from subject matter experts
Link	Relations among entities	Causal link from one belief to another	Causal link from one belief to another
Node Parameter	None	Predefined baseline probability of belief's becoming true	Expert's baseline probability of belief's becoming true
Link Parameter	Edge weight showing the strength of the relation	Predefined promotion and inhibition	Expert's promotion and inhibition

3. Dataset

Throughout this paper, we use a dataset collected from 1998 US Embassy bombing incident in Kenya. The dataset is a meta-network of a terrorist organization. This dataset is initially extracted from a network text analysis on open-source documents, but later, the dataset went through corrections by human analysts. As our framework starts with a meta-network, the initial input dataset is a collection of terrorists, information and resources for the bombing, and related tasks. Figure 1 is the visualization of the meta-network of the Kenya case. The basic statistics of this network is listed in Table 2. For each of the sub-networks, there is an interpretation for the links. For instance, the link in a social network represents that two terrorists interacted or communicated with each other, and the link in a task assignment network shows that the terrorist was assigned to completion of the linked task.

Table 2: The meta-matrix of the dataset, a terrorist group responsible for 1998 US embassy bombing in Kenya, The numbers in the cells represent the densities of the sub-networks.

	Terrorist	Expertise	Resource	Task

Terrorist (16 terrorists)	Social Network (0.141)	Information Distribution Network (0.078)	Resource Distribution Network (0.093)	Task Assignment Network (0.134)
Expertise (8 knowledge nodes)		Not used	Not used	Required Expertise Network (0.048)
Resource (8 resources)			Not used	Required Resource Network (0.076)
Task (13 tasks)				Task Precedence Network (0.121)

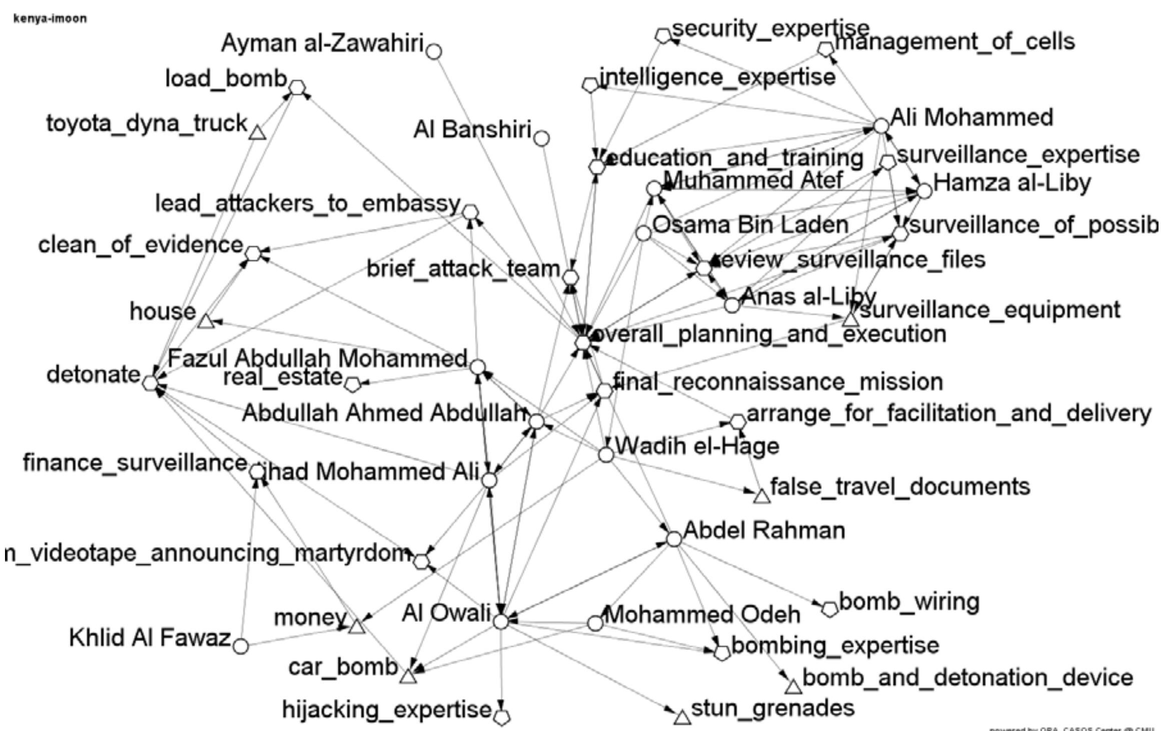

Figure 1: The visualization of the meta-matrix of the terrorist group responsible for the 1988 US embassy bombing in Kenya

4. Method - Generating an influence network from a social network

We generate an influence network explaining the likelihood of a task completion from a social network. Thus, inputs for the generation are 1) a social network, 2) a target task to be analyzed in the network and 3) parameters for the generated network. Political, Military, Economic, Information and Infrastructure, or PMESII, aspects are the elements of assessing a situation. We similarly identify six factors contributing to a task completion. The six factors are 1) prior task completion, 2) task importance, 3) task complexity, 4) personnel assignment, 5) accessible expertise and 6) available resources. The below sections explain how we extracted each of the task completion factors and turn them into a node in the influence network.

4.1. Overall structure of a generated influence network

We describe the overall structure and how the accompanying parameters are determined. The structure of a generated influence network is explained in two steps. First, the skeleton of the influence network is from the task network of a particular final task. In the social network, there is a task network specifying the prior and the next tasks of a certain task. If an analyst selects a task to be analyzed, we infer a sub-network that only selects the tasks related to the completion of the final task and create a task network for it. That becomes the skeleton of the influence network. After that, we assess the likelihood of success for each task by adding the above six factors as influence network nodes. This becomes the flesh of influence network modeling the success of the each task. With these two parts, we can propagate estimation on the success likelihood of individual task throughout the influence network with the skeleton of task network, Figure 2.

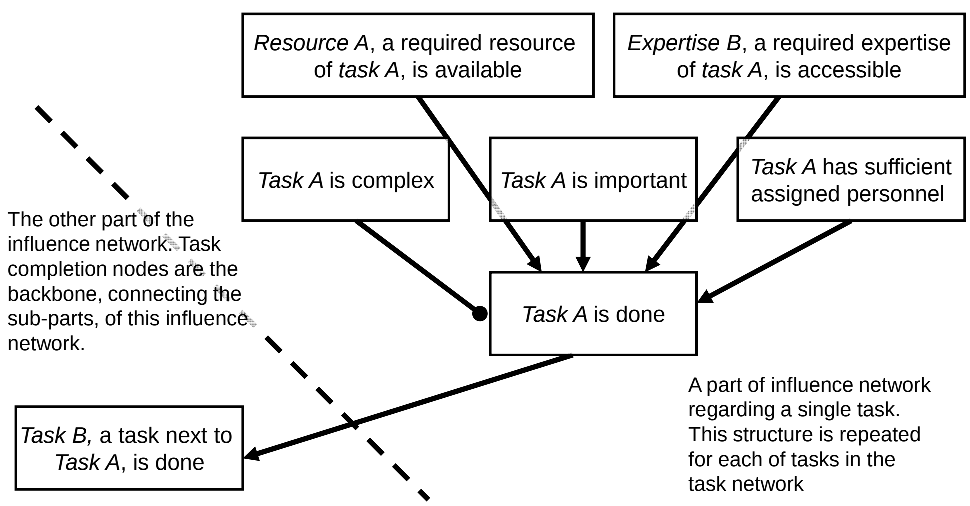

Figure 2: A simple diagram displaying how a generated influence network is structured. The skeleton of the network is the task network of a target task. Then, each task has five factors related only to the task.

While we set the topology of the influence network as above, we supply a set of heuristics determining the accompanying parameters for the network. This heuristic contains our assessment criteria on how to organize personnel, resources and expertise to successfully execute a task. Each of the factors and the heuristics are explained in Table 3 and the below sections.

Table 3: A summary of assigning a marginal probability for each of identified influence network node. The baseline probability assignment is uniform across the nodes from the

same factor, Two used network metrics, degree and betweenness centralities, are defined in Freeman (1979)

Factor name	Marginal probability	Organizational structure assessment
Task A is done (Task network)	Medium:0.5	Always Medium level baseline probability. The completion of a task has 50% of chance if there is no external influence.
Task A is Complex (task complexity)	Very Low: 0, Low: 0.25, Medium:0.5, High: 0.7, Very High: 0.8	- Very Low marginal probability if 1 person and 0 resource/expertise required - Low marginal probability if 2 persons and 1 resource/expertise required - Medium marginal probability if 3 persons and 3 resources/expertise required - High marginal probability if 6 persons and 7 resources/expertise required - Very High marginal probability for the rest of cases
Task A is important (task importance)	Very Low: 0, Low: 0.25, Medium:0.5, High: 0.7, Very High: 0.8	- Very Low marginal probability if 0 degree or 0 betweenness centrality - Low marginal probability if 0 - 0.25 degree or 0 - 0.25 betweenness centrality - Medium marginal probability if 0.25 - 0.5 degree or 0.25 - 0.5 betweenness centrality - High marginal probability if 0.5 - 0.75 degree or 0.5 - 0.75 betweenness centrality - Very High marginal probability for the rest of cases
Task A has sufficient assigned personnel (personnel sufficiency)	Very Low: 0, Low: 0.25, Medium:0.5, High: 0.7	- Very Low marginal probability if 0% of required resources and expertise are covered by the assigned personnel - Low marginal probability if 50% of required resources and expertise are covered by the assigned personnel - Medium marginal probability if 75% of required resources and expertise are covered by the assigned personnel - High marginal probability if 100% of required resources and expertise are covered by the assigned personnel
Resource A, a required resource of task A, is available (available resources)	Very Low: 0.25, Low: 0.5, Medium:0.75,	- Very Low marginal probability if the task has 0 assigned personnel with the required resource - Low marginal probability if the task has 1 assigned personnel with the required resource - Medium marginal probability if the task has 2 or more assigned personnel with the required resource

While Table 2 specifies how we determine the marginal or baseline probabilities of the influence network nodes, each of the influence network links requires two parameters: promotion and inhibition weights. The promotion weight is the strength of the influence toward the destination influence network node when the start node is true. The inhibition is the influence strength to the destination node when the start node is false. Throughout this paper, we use 0.5 for promotion and -0.5 for inhibition weights. These values are selected because we want to balance the causal strengths regardless of the success of the parent nodes. These weights can change as human analysts' qualitative assessment of a target situation. If human analysts feel that the failure of a task facilitates the failure of the subsequent tasks more than the task success promotes the subsequent task successes, they should decrease the promotion weight and increase the inhibition weight.

4.2. Task network

Unlike the other five factors, the effect of prior task completion propagates to the child tasks throughout an influence network. For instance, if task A is a prior task of task B, and task B is that of task C, the likelihood of task A affects that of task C. This is different from the other factors, i.e. task complexity of a certain task contributes the task's completion likelihood in a negative way, but this contribution is limited to that task. This propagation relation can be extracted from the task network in a social network. The social network that we use is in meta-network format, which specifies the task-to-task network. If the task network has directionality, we can see the task flow from the leaf to a certain task. For example, Figure 3 shows an extracted task network, a task network for detonate task. Because prior task completion is the only factor with propagation attribute, we build up an influence network from this task network for a specific task. Then, we can add the other five factors to each of the task in the influence network already.

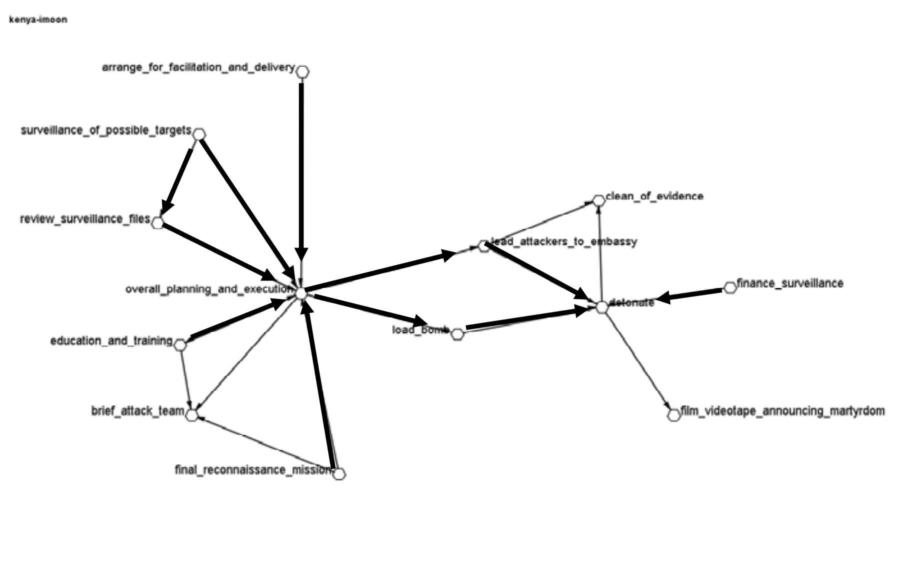

Figure 3. the task network from the Kenya case. The sub-task network of *detonate* is highlighted.

4.3. Task importance

A task is more likely to be executed successfully if the task is considered to be important. Therefore, as the task importance of a specific task goes up, the task completion likelihood increases and task importance has a promoting influence on task completion. Then, the question is how to measure the importance of each task in the task network. We gauge the importance based on the number of prior and following task in the task network. If a task has many prior or next tasks, the task is important. This factor can be measured by the degree centrality of a task in a task network. Also, if a task is on many critical paths among two tasks in the task network, the task is important. This is captured by measuring the betweenness centrality of a task. For instance, in the task dependency network of Figure 4, overall planning and execution is a task with 0.4167 degree centrality and 0.2652 betweenness centrality, so the task is considered to be important with 0.7 marginal probability. Each task node in the influence net has the task importance factor node as a parent in the influence network, and the importance node probability is calculated from the heuristics as described in the previous section.

4.4. Task complexity

A task is less likely to be performed if the task has high complexity. In a social network, 'a task is complex' means that the task requires many personnel involvements and different types of entities. Thus, we measure the task complexity factor with the number of assigned agents and the number of required resources and expertise. For example, Figure 4 shows two tasks, detonate and review surveillance files. The former has one associated agent, jihad mohammd ali, and two required resources, toyota dyna truck and car bomb. The latter has five related agents, anas al liby, hamza al liby, muhammed atef, ali Mohamed, and osama bin laden, and one required expertise, surveillance expertise. Because the latter requires much more personnel and similar number of resources and expertise, the latter has a higher baseline probability, 0.5, for the task complexity factor than the former, 0.25. This task complexity becomes a node in the influence network and is linked to the task node from section 3.1.1.

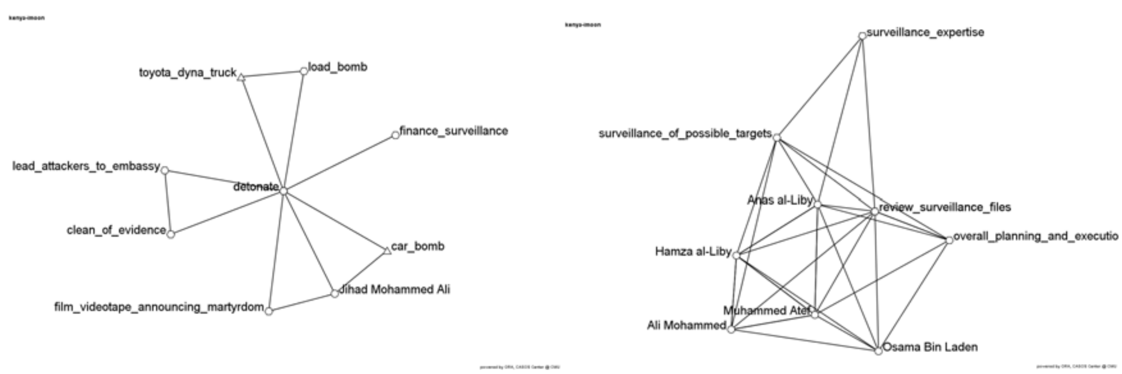

Figure 4: (top) A sub-network including nodes in one social link distance from detonate, (bottom) a sub-network from review surveillance files

4.5. Personnel sufficiency

A higher personnel sufficiency is a key element in the task completion. However, we have seen that only providing an agent without any proper resources or expertise is not enough. Therefore, when we count the personnel sufficiency, we see not only the number of agents, but also whether the agent has a required resources or expertise. For instance, as Figure 5, overall planning and execution task has no required resources and expertise, and it has eight assigned agents. Therefore, its personnel sufficiency level, 0.7, is high. On the other hand, detonate task requires two resources. However, it has only one agent with one required resource, so the personnel sufficiency of detonate task is low as 0.25.

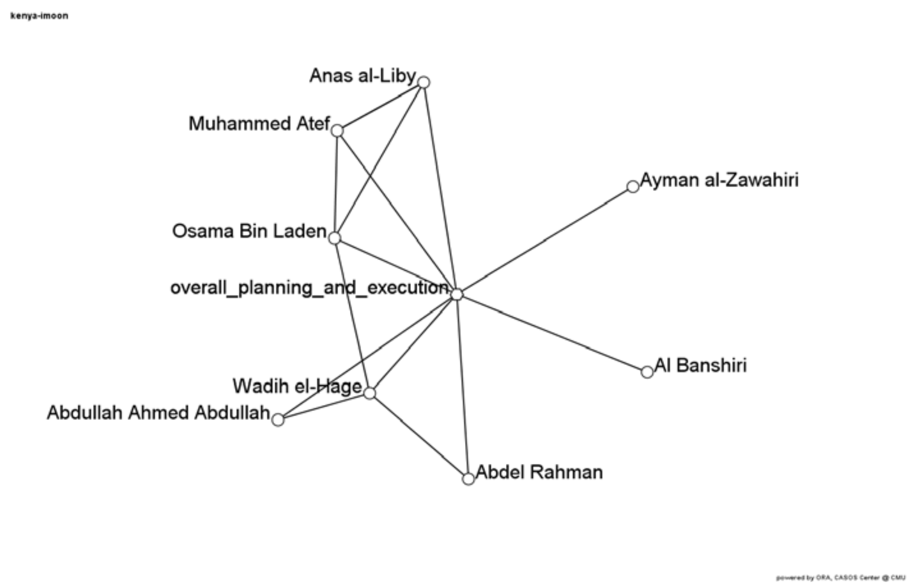

Figure 5: The sub-network of nodes within one social link distance from overall planning and execution. When we limit the types to agents, resources and expertise, there is no required resources and eight assigned agents.

4.6. Accessible expertise and available resources

Finally, providing required expertise and resources of a task to assigned agents is an important factor in task completion. In detonate task, Figure 4, car bomb is provided to Jihad Mohammed Ali, but Toyota Dyna Truck is not assigned to anyone doing the task. Therefore, an influence node, car bomb is available, has a higher baseline probability, 0.25, compared to that of Toyota Dyna Truck is available, 0. This idea applies to expertise as well as resources.

5. Result

We apply the above influence network generation algorithm to the Kenya case. First, we perform a sensitivity analysis of the parameters we assigned. Next, we see the completion likelihood of a key task, detonate.

5.1. Baseline probability sensitivity analysis

We formulate our problem in a very basic form. A final task node has five factors (leaf nodes)—*resource availability, personnel assignment, task importance, task complexity*, and *expertise*—that affect the marginal probability. These factors have different baseline probability assignments in *Organization Risk Analyzer* (or ORA, Carley et al., 2007) based on our preconceived notion of how these factors may weigh differently. Now we would like to investigate how sensitive a final task node's marginal probability is given different baseline probability assignments.

For simple calculation we assumed the independence of five factors {*resource availability, personnel assignment, task importance, task complexity,* and *expertise*}. In reality we presume there are some correlations among these factors' probability *very low/ low/ medium/ high/ very high* levels. Each factor's baseline probability drawing was based on the "U.S. Embassy bombing in Kenya" example. There *resource availability* has a uniform distribution; *personal assignment* has a right-skewed shape (as shown in Fig 6-left); *task importance* is fixed at 'medium' probability; *task complexity* has a left-skewed shape (as shown in Fig. 6-right); and *expertise* has a close to half-and-half chance of 'medium' and 'high' probability and 1% chance of 'low' probability.

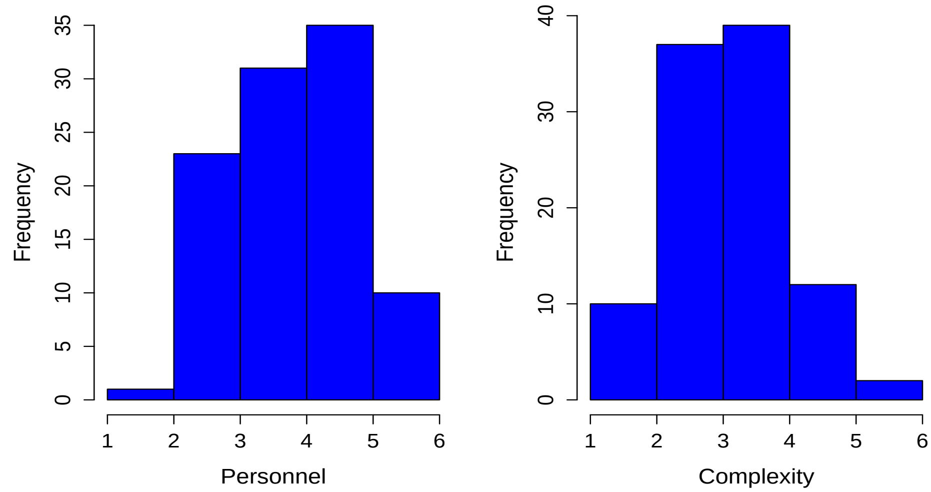

Figure 6: *Personnel assignment* and *complexity* factors' probability assignment distributional shapes from 100 independently simulated samples. 1 corresponds to 'very low' probability and the number increases up to 5 for 'very high' probability.

The baseline probabilities in ORA tend to give an upward shift of the final marginal probability by 5% up to 30% in comparison to the set of baseline probabilities that are fixed across factors where we assign the same "standardized values" {0.1, 0.3, 0.5, 0.7, and 0.9} to every factor's *very low/ low/ medium/ high/ very high* baseline probabilities so that the factors are indistinguishable by weight. See Figure 8 where we plot the discrepancy between the two baseline probability assignments' 1,000 simulation results for each of four different settings. The top-left shows the histogram of marginal probabilities from the "U.S. Embassy bombing in Kenya" example; the top-right shows the simulation discrepancy results where *resource availability, personnel assignment, task complexity* factors' probability assignments were drawn from independent uniform distributions; the bottom-left differs from the top-right case only in the *personnel assignment* with the right-skewed distribution as observed from the "U.S. Embassy bombing in Kenya" example; and the bottom-right is changing from the top-right case only in the *task complexity* to a left-skewed distribution. (See Figure 6 for the two non-uniform factors' distributions.) In our simulation, about 80% of the ORA assigned marginal probabilities report 10% to 20% more successful outcome/execution of the final task regardless of the four experimental settings we distinguished. That is, the current ORA baseline probability assignment interprets the operational process of adversaries as a more efficiently working field by assigning highly positive weight on the three factors that were not tweaked: *resource availability*, *task importance*, and *expertise*. The current ORA probability assignment setting is in Table 3.

Another phenomenon we see from Figure 7 is that two pairs of histograms on the diagonals display very similar shapes. We can see that keeping *personnel assignment* factor uniform does not change much the pair-wise outcome from the Kenyan example, as well as keeping three factors—*resource availability, personnel assignment, task complexity*—uniform renders similar outcome as changing the *personnel assignment* as right-tail heavy since we already have marginals' overestimates from the original pair-wise comparison example. It is like a convolution with a changing distribution-shape kernel.

Though *expertise* is biased toward a "medium-to-high" probability level, its effect on the marginal probability was minimal. When all other factors were input as 0.5 (baseline probability), the final marginal probability followed the same distribution as the *expertise* probability level distribution. About half the marginals remained at 0.5 and the other half was about 0.55.

When all the factors were combined and effective, *personnel assignment* influence more than the *task complexity* where we see more spread out outcome of marginal probabilities. We need to further investigate how much other factors have played a role in this marginal probability distribution's thicker tails.

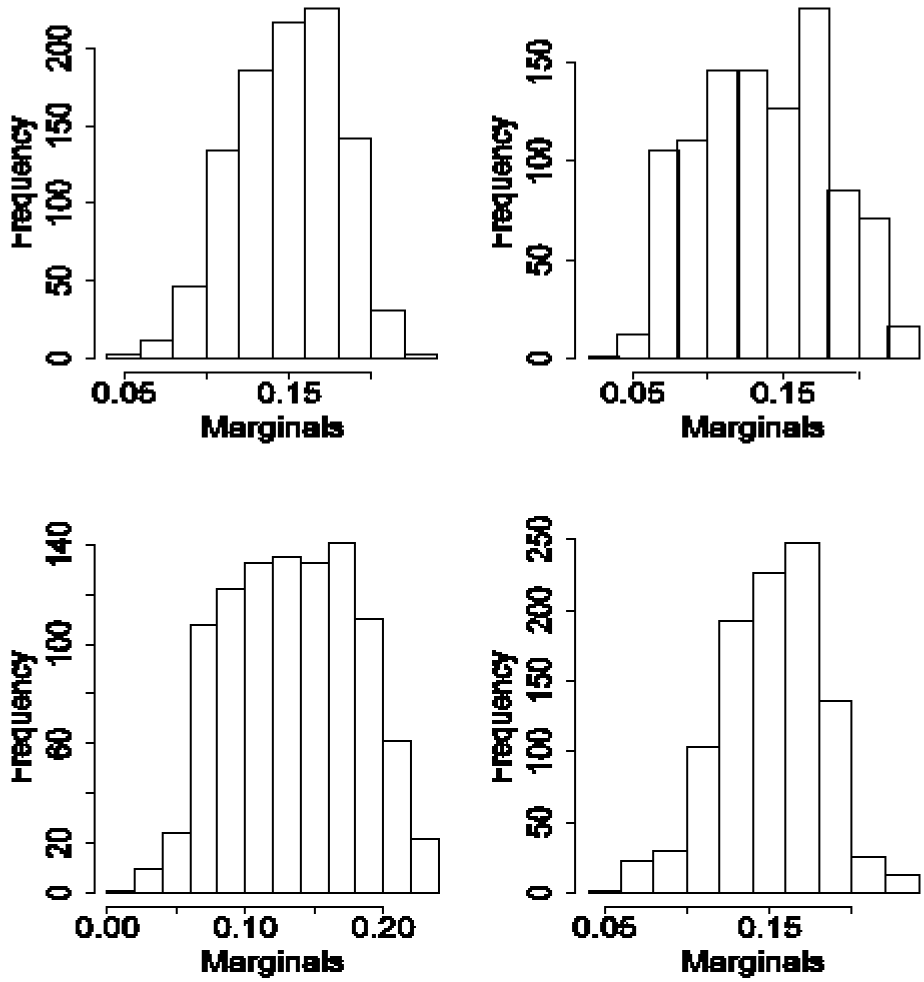

Figure 7: Histograms of the differences (ORA default minus the "Standardized" baseline probability assignments) in the final task node's marginal probability for four different experiments.

We draw another comparison between the ORA default version and the "standardized" version where we find that the spread or the interval of the probability assignments matter. With a fixed 0.2 interval in our standardized version, we have the marginal probability piling up around 10 values, which implies that the size of the interval tells us a more clear-cut estimation of the marginal probability. On the other hand, in the current setting of ORA where we have the baseline probability assignment intervals varying across factors, the final task node's marginal probability has a distribution rather smooth and it mostly ranges over high probability (between 0.5 and 0.8). See Figure 8.

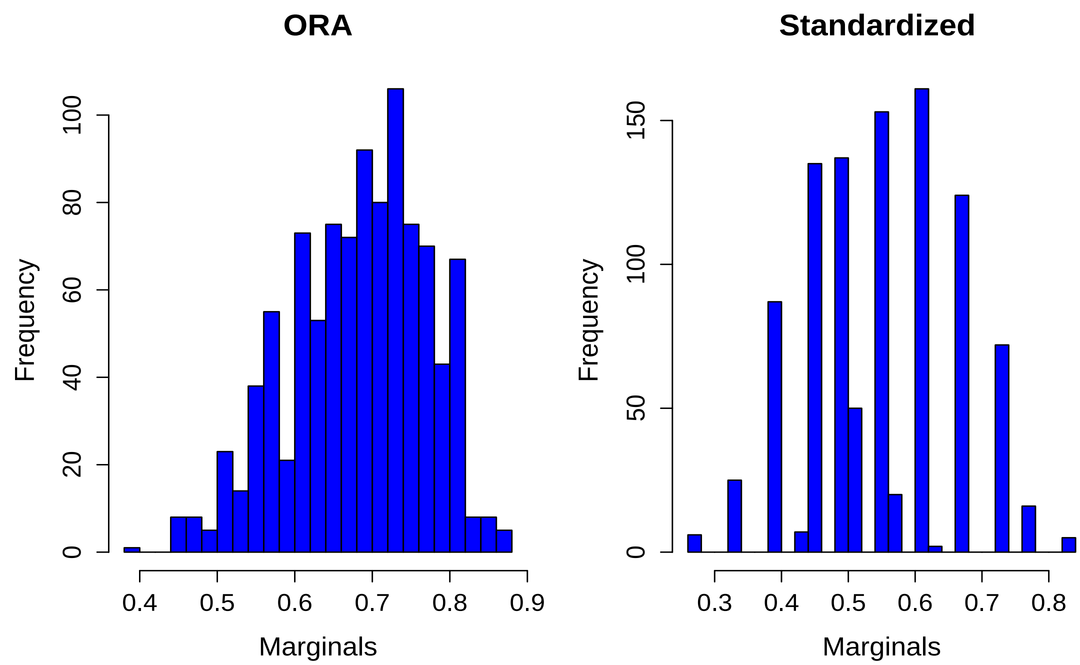

Figure 8: Histograms of marginal probabilities where the left panel shows the marginal probabilities based on the ORA-default baseline probabilities and the right panel shows in case of using the standardized baseline probabilities.

5.2. Assessment on completion likelihood of a key task

Figure 9 is the visualization of a generated influence network using Pythia (Wagenhals and Levis, 2007). Pythia is a computer program that subject matter experts use to design and evaluate an influence network. This influence network is of the Kenya US Embassy bombing case, and its target task is the *detonation*. While ORA generates the influence network, Pythia evaluates the marginal probability for each of the influence network node. Because the factors regarding a single task are leaf nodes, we do not need marginal probability of those. However, the nodes in the skeleton of this influence network are the tasks that we want to know of their success likelihood, whose values change from its baseline probability depending on the leaf nodes.

Table 4 displays the evaluation result of the task nodes. The final task detonation, in the above influence network has 0.36 likelihood of success, but in reality the organization was successful in executing the task. When we examine the task network of detonation, we see the major contribution of such a low probability comes from the low likelihood of Load bomb. Detonationg had three required tasks: Load bomb, Lead attackers of the embassy, and Finance surveillance. While the first task had the success likelihood at 0.31, the latter two had the likelihood at 0.54 and 0.34 respectively. Therefore, to increase the success likelihood of the final task above 36%, this organization should strengthen the support from Load bomb execution. Friendly forces should undermine Lead attackers of the embassy and Finance surveillance since those two tasks still had high likelihood. This is just a brief analysis using the automatic network generation function of ORA and the evaluation function of Pythia.

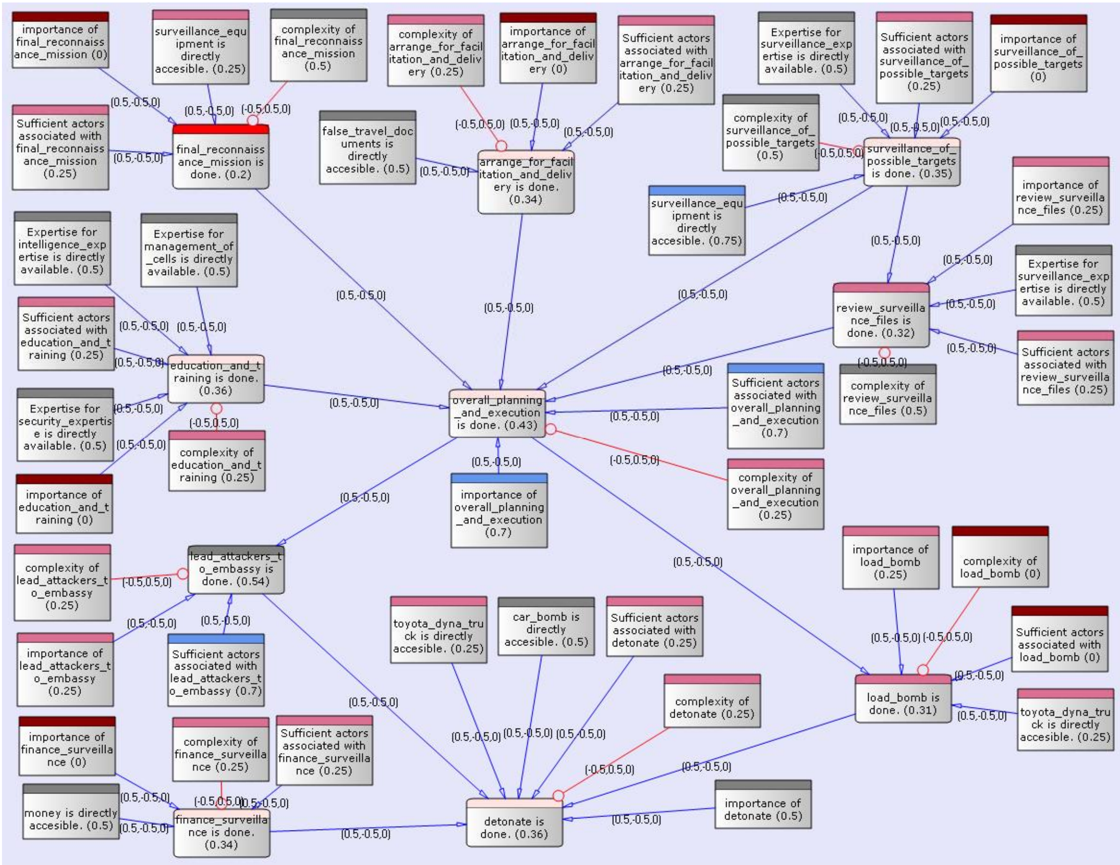

Figure 9: The visualization of an influence network from Pythia:
The analysis target task is detonation.

Table 4. Evaluation result from the generated influence network for detonation task

Task Name	Overall planning and execution	Surveillance of possible targets	Detonate	Load bomb
Marginal Prob.	0.43	0.35	0.36	0.31
Task Name	Education and training	Finance surveillance	Review surveillance files	Final reconnaissance mission
Marginal Prob.	0.36	0.34	0.32	0.20
Task Name	Lead attackers to embassy	Arrange for facilitation and delivery		
Marginal Prob.	0.54	0.34		

6. Conclusion

We introduced the automatic generation of an influence network from a social network. Also, we provided an illustrative example of the usage by applying it to the Keyna case dataset. The brief analysis reveals that the adversaries are successful while the key final task, detonate, has 36% of marginal probability, or task completion likelihood. To reduce this number, the model suggests that the friendly forces should decrease the likelihood of Lead attackers to embassy, Finance surveillance or both. It is already identified that the adversaries have such a small task completion likelihood of Load bomb, which is one of the prior tasks of detonate.

There are many other analyses that we can integrate into this framework. First, the course of action (COA) generator of Pythia can be directly applicable. The COA generator will tell which sector or factor to work on to decrease the likelihood of the final task. Also, we can apply the strategic intervention concept from social network to this model. We remove a set of nodes, regenerate an influence network from the modified social network, evaluate the likelihood again and compare its drop. If the drop is huge, then the set of removed nodes are the nodes that friendly forces should get rid of. This integration will contribute to the counterterrorism analysis field by 1) helping the analysts trying to create an influence network from a scratch by providing a basic influence network that they can work on and 2) facilitating the creation of unified analysis framework that can be broadly used in the intelligence and the military planning fields.

Acknowledgement

This work was supported in part by the Air Force Office of Sponsored Research (MURI: Cultural Modeling of the Adversary, 600322) for research in the area of dynamic network analysis. Additional support was provided by CASOS - the center for Computational Analysis of Social and Organizational Systems at Carnegie Mellon University. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the AFOSR.

References

Argote, L. and Ingram, P. (2000) Knowledge Transfer: A Basis for Competitive Advantage in Firms, Organizational Behavior and Human Decision Processes, Vol. 82, No. 1, pp. 150-169

Becker, B. and Gerhart, B. (1996) The Impact of Human Resource Management on Organizational Performance: Progress and Prospects, The Academy of Management Journal, Vol. 39, No. 4, pp. 779-801

Burke, J. (2004). Al Qaeda: The True Story of Radical Islam. London, Penguin.

Carley, K. M. (2006) Destabilization of covert networks, Computational & Mathematical Organization Theory, Vol. 12, Issue 1, pp 51-66

Carley, K. M., Columbus, D., DeReno, M., Reminga, J. and Moon, I. (2007) ORA User's Guide 2007, Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISRI-07-115.

Campbell, D. K. (1991) Goal Levels, Complex Tasks, and Strategy Development: A Review and Analysis, Human Performance, Vol. 4, No. 1, pp. 1-31

Defense Advanced Research Projects Agency (DARPA) (2005) Integrated Battlefield Command (IBC) Program Solicitation Briefing. [Online] Available: http://www.darpa.mil/sto/solicitations/IBC/.

Eisner, H. (1962) A Generalized Network Approach to the Planning and Scheduling of a Research Project, Opera-tions Research, Vol. 10, No. 1, pp. 115-125

Forsyth, D. R. and Schlenker, B. R. (1977) Attributing the causes of group performance: Effects of performance quality, task importance, and future testing, Journal of Personality, Vol. 45, Iss. 2, pp. 220-236

Freeman, L. C. (1979) Centrality in Social Networks: Conceptual Clarification, Social Networks 1.3: 215-239.

Hoffman, B. (1998) Inside Terrorism. London: St. Andrew's University Press

Hudson, L. D., Ware, B. S., Mahoney, S. M. and Laskey, K. B. (2001) An application of Bayesian networks to anti-terrorism risk management for military planners, Technical Report, Department of Systems Engineering and Operations Research, George Mason University

Hufbauer, G. C., Schott, J. J. and Oegg, B. (2001) Using Sanctions to Fight Terrorism. International Economics Policy Briefs, Vol. 11, Num. 1, pp 1-19

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

Pentland, B. T. and Feldman, M. S. (2007) Narrative Networks: Patterns of Technology and Organization, Organization Science, Vol. 18, Num. 5, pp 781-795

Rosen, J. A., and Smith, W. L. (1996) Influence Net Modeling with Causal Strengths: An Evolutionary Approach, Proceedings of the Command and Control Research and Technology Symposium, Naval Post Graduate School, Monterey CA, Jun.

Silverman, B. G., Bharathy, G., Nye, B. and Eidelson, R. J. (2007) Modeling factions for "effects based operations": part I-leaders and followers. Computational & Mathematical Organization Theory, Vol. 13, Num. 4, pp 379-406

Vego, M. N. (2006) Effects-Based Operations: A Critique. Joint Forces Quarterly, Issue 41, pp 51-67

Wagenhals, L. W. and Levis, A. H. (2007) Course of Action Analysis in a Cultural Landscape Using Influence Nets. IEEE Symposium on Computational Intelligence in Security and Defense Applications, Apr 1-5, pp 116-123