FROM THE USMA-ARI WORKSHOP

Social Network Monitoring of Al-Qaeda

Ian McCulloh, Kathleen M. Carley and Matthew Webb

ABSTRACT

Social network monitoring is the application of statistical process control charts to changing social network measures over time. Quality engineers use control charts to detect slight changes in industrial manufacturing processes. Once detected, the quality engineer will identify a maximum likely change point, when the process began to change, and search for the specific cause of the change. These tools allow quality engineers to quickly identify changes before they cause significant financial loss to the manufacturing company. In the same manner, analysts can use control charts to detect slight changes in dynamic social network measures.

A cumulative sum (CUSUM) control chart is applied to a dynamic social network data set of the Al-Qaeda terrorist organization. The data set ranges from 1988 to 2004. The CUSUM identifies a shift in several network measures in the Al-Qaeda network between 2000 and 2001. The CUSUM most likely change point for all measures is 1997. This example suggests that if analysts were to use social network monitoring to monitor terrorist networks, dangerous shifts in the network might be detected before they become a problem. Furthermore, the specific cause of change could be identified, allowing analysts to exploit positive changes in a terrorist organization and mitigate negative changes.

INTRODUCTION

Social network analysis (SNA) is the mathematical methodology of quantifying relationships between individual people or organizations. SNA considers individuals or groups as nodes in a graph, while relationships between nodes are graph edges. There can exist many different types of relationships, such as communication, finance, religion, or nationality. This methodology offers a wealth of potential tools for military intelligence and the war on terror.

The Center for Computational Analysis of Social and Organizational Systems (CASOS) maintains social network data on the Al-Qaeda terrorist network, developed under a research grant from the Office of Naval Research (ONR). This data includes many different relationships to include communication, financial, physical, etc. The data set begins with intelligence collected in 1988 and includes consecutive years through 2004. Using this data and SNA methods, analysts are able to calculate and quantify the most influential terrorists in the network, the most knowledgeable, individuals that connect separate subsections within the group, and much more. While it is important to understand terrorist organizations from an SNA perspective, it does not necessarily identify critical changes in social network structure over time.

Statistical process control charts may be useful in monitoring social networks for important changes over time. Control charts are used by quality engineers in industry to monitor manufacturing processes for changes to important quality characteristics. The quality engineer records observations of a specified quality measure and calculates an appropriate statistic. He then compares the statistic to a control limit. If the statistic exceeds the control limit, the chart is said to "signal" that there may have been a change to the quality characteristic that is being monitored. The quality engineer will then inspect the process to see if it is out of calibration, before the process continues to produce product that is outside quality specifications. He then corrects the process if necessary; and begins to search for the specific cause of the signal. Some control chart algorithms offer an estimate of when the process fell "out-of-control". This saves time in identifying the specific cause of the signal. These same control chart algorithms can be applied to SNA measures observed on networks collected over time. Instead of observing quality characteristics, however, normally distributed network measures are used.

A cumulative sum (CUSUM) control chart is a commonly used statistic in quality engineering that offers an estimate of when the observed process changes. The CUSUM statistic is therefore applied to several normally distributed measures from the ONR Al-Qaeda SNA data from 1988-2004. Using an arbitrary control limit, the control chart is able to successfully predict changes to the Al-Qaeda network prior to their terrorist attacks of September 11th.

ISBN 978-1-934808-00-9

CUSUM STATISTIC

The CUSUM control chart is a widely used control chart derived from the sequential probability ratio test (SPRT) (Page, 1961). The SPRT was derived from the Neyman and Pearson (1933) most powerful test for a simple hypothesis. Neyman and Pearson's test statistic is

$$\Lambda_t = \frac{\prod_{i=1}^{t} f(x_i;\mu_1)}{\prod_{i=1}^{t} f(x_i;\mu_0)} \quad . \qquad (2.1)$$

Neyman and Pearson showed that the most powerful test of H_o against H_1 is obtained by rejecting H_o if $\Lambda_t \geq K$ and concluding in favor of H_o if $\Lambda_t < K$, where K is determined by the level of significance, α. The level of significance is the probability that H_o is rejected when it is true.

Wald (1947) demonstrated that the Neyman and Pearson hypothesis testing method could be applied sequentially and could significantly reduce the number of samples required to reach a conclusion. Wald's sequential probability ratio test (SPRT) compares Λ_t to two constants A and B where $0 < B < A < \infty$. Observations are collected and examined one-at-a-time. After the t^{th} observation there are three possible outcomes. If $\Lambda_t < B$, then the test concludes in favor of H_o. If $\Lambda_t > A$, then H_o is rejected in favor of H_1. If $B \leq \Lambda_t \leq A$ then sampling will continue with observation $t + 1$.

The SPRT can be used to test H_o: $\mu = \mu_o$ against H_1: $\mu = \mu_1$ for normal means. Without loss of generality we will assume that $\mu_1 > \mu_o$. Having observed t observations, the SPR is

$$\Lambda_t = \frac{\prod_{i=1}^{t}\left(\frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{1}{2}\left(\frac{x_i-\mu_1}{\sigma}\right)^2\right)\right)}{\prod_{i=1}^{t}\left(\frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{1}{2}\left(\frac{x_i-\mu_0}{\sigma}\right)^2\right)\right)} \quad .$$

This can be reduced algebraically to

$$\Lambda_t = \exp\left(\left(\frac{\mu_1-\mu_0}{\sigma^2}\right)\sum_{i=1}^{t} x_i + t\left(\frac{\mu_0^2-\mu_1^2}{2\sigma^2}\right)\right) \quad . \qquad (2.2)$$

The sequential probability ratio, Λ_t, is compared to appropriate constants A and B as each new observation t is formed. Following observation t, the test concludes in favor of H_o if $\Lambda_t < B$. If $\Lambda_t > A$, then the test concludes in favor of H_1. If $B \leq \Lambda_t \leq A$, then sample $t + 1$ is obtained and a revised Λ_{t+1} is computed. This procedure continues until either $\Lambda_t < B$ or $\Lambda_t > A$.

In a Social Network Monitoring (SNM) application of the SPRT, μ_o is the mean of a social network measure and μ_1 is the mean after a change in the network. Since one would never conclude in favor of H_o that the network is unchanged and stop all sampling, the procedure continues until it signals that there is a change in the network. This implementation of the SPRT procedure leads to the CUSUM control chart.

The CUSUM control chart is based on cumulative sums of a network measure over time and is derived from the sequential probability ratio test (SPRT). In a control chart application of the SPRT, one would continue to monitor the network until $\Lambda_t > A$ when the procedure signals that there is a change in the network. The SPRT leads to the following expression for detecting an increase in the mean of a normally distributed network measure. The procedure would signal when

$$\Lambda_t = \exp\left(\left(\frac{\mu_1-\mu_0}{\sigma^2}\right)\sum_{i=1}^{t} x_i + t\left(\frac{\mu_0^2-\mu_1^2}{2\sigma^2}\right)\right) > A \quad . \qquad (2.3)$$

This expression can be simplified by taking the natural logarithm of both sides of the inequality,

$$\left(\frac{\mu_1-\mu_0}{\sigma^2}\right)\sum_{i=1}^{t} x_i + t\left(\frac{\mu_0^2-\mu_1^2}{2\sigma^2}\right) > \log A$$

This decision rule can be algebraically reduced to

$$\sum_{i=1}^{t} x_i - t\left(\frac{\mu_0+\mu_1}{2}\right) > A'$$

, where

$$A' = \left(\frac{\sigma^2}{\mu_1-\mu_0}\right)\log A \quad .$$

By allowing $\mu_1 = \mu_o + \delta\sigma_x$, the procedure signals when

$$\sum_{i=1}^{t} x_i - t\left(\frac{\mu_0+(\mu_0+\delta)}{2}\right) = \sum_{i=1}^{t}\left(x_i - \mu_0 - \frac{\delta}{2}\right) > A'$$

where δ is the standardized difference in the network measure under H_o and H_1. This decision rule can then be further simplified by using the cumulative statistic

$$C_t = \sum_{i=1}^{t}(Z_i - k)$$

where $Z_i = (x_i - \mu_o)/\sigma_x$, and $k = \delta / 2$. The common choice of k in quality applications is 0.5, which corresponds to a standardized magnitude of change in

mean of $\delta = 1$. Thus, observations are examined sequentially until $C_t > A'$.

The CUSUM sequentially compares the statistic C_t against a control limit A' until $C_t > A'$. Since one is not interested in concluding that the network is unchanged, the cumulative statistic is

$$C_t^+ = \max\{0, \quad Z_t - k + C_{t-1}^+\} \qquad (2.4)$$

If this rule was not implemented the control chart would require more subgroups to signal if $C_t < 0$. The statistic C_t^- is compared to a constant, h^-. If $C_t^- > h^-$, then the control chart signals that an increase in a network measure has occurred.

A required assumption for the derivation of the CUSUM statistic is that the network measure under observation is normally distributed. There is still much work to be done in the area of classifying the probability space of a social network. However, the central limit theorem does allow us to understand the distribution of a sample average of 30 or more observations. Therefore, network measures that are averaged over 30 or more nodes will have a normally distributed measure. Some of the measures that have been investigated include the Average Betweenness, Average Closeness, Average Degree, and Average Eigenvector Centrality.

For $\delta < 0$, the SPRT similarly leads to the CUSUM procedure for detecting a decrease in a network measure. In this case,

$$C_t^- = \max\{0, \quad -Z_t - k + C_{t-1}^+\}$$

is compared to a constant, h^-. If $C_t^- > h^-$, then the control chart signals that a decrease in a network measure has occurred.

To monitor for both increases and decreases in the mean, two one-sided control charts are employed. One chart is used for monitoring for increases in the mean of a network measure and the other is used for detecting decreases in mean. If the process remains in-control, C_t^+ will fluctuate around zero. If there is an increase in the mean of a network measure, C_t^+ will tend to increase. Conversely, if there is a decrease, then C_t^- will tend to increase. When $C_t^+ > h^+$ or $C_t^- > h^-$, the two one-sided CUSUM control chart scheme signals that the process is out-of-control.

The CUSUM control chart's ability to detect changes has been extensively investigated in the literature. Lorden (1971) introduced a minimax criteria that minimizes the average number of observations to detect a change, subject to a given probability of false alarms. He also proposed the use of a maximum likelihood approach to rapidly detect changes in a process. Lorden's approach does better than the CUSUM at detecting a wide range of changes in the mean of a process, but does not completely outperform the CUSUM for all potential changes in the process mean.

Moustakides (1986) proved that Page's CUSUM control chart will detect a specific standardized step change in the mean of a measure with fewer subgroups than any other statistical test. The specific step change in mean is a standardized change of $\delta = 2k$, where k is the control chart parameter in Equation 2.4. Therefore, the CUSUM control chart is the best chart to use for detecting a standardized change in mean of $\delta = 2k$. However, other control charts may detect other changes in the mean with fewer subgroups. Several attempts to improve upon the CUSUM control chart have been investigated in the literature. However, the CUSUM is used to demonstrate the general applicability of control charts in dynamic social network analysis, due to its simplicity and versatility.

RESULTS

Social network measures were plotted for number of agents, average degree, average betweenness, average closeness, average eigenvector centrality, and density. Each of these network measures were increasing from 1988 until 1994. The measures then leveled off. There are many possible reasons for this burn-in period, the least of which is the quality of intelligence gathering on Al-Qaeda. For this reason, the average measure and standard deviation were calculated over five years beginning in 1994. The CUSUM control chart was used to monitor the five measures above from 1994 to 2004. Figure 1 displays the plot of the social network measure for the average closeness of members in the Al-Qaeda network. The reference value, k, and the control limit, h, were arbitrarily set at 0.5 and 4 respectively for all of the social network control charts. Figure 2 shows the CUSUM statistic for the average closeness that is plotted in Figure 1. It can be seen that the CUSUM statistic in Figure 2 is a more dramatic indication of network change than simply monitoring the network measure in Figure 1. This is a result of the CUSUM statistic taking into account previous observations of the network. A single observation of a network measure that is slightly higher than normal may not indicate a change in the network. Multiple observations that are slightly higher than normal, however, may indicate a shift in the mean of the measure.

Recall that the CUSUM will either detect increases or decreases in a measure, but not both. Therefore, two control charts must be run for each social network measure being monitored. One chart

ISBN 978-1-934808-00-9

Figure 1. Control Chart for the Average Closeness of Al-Qaeda Members

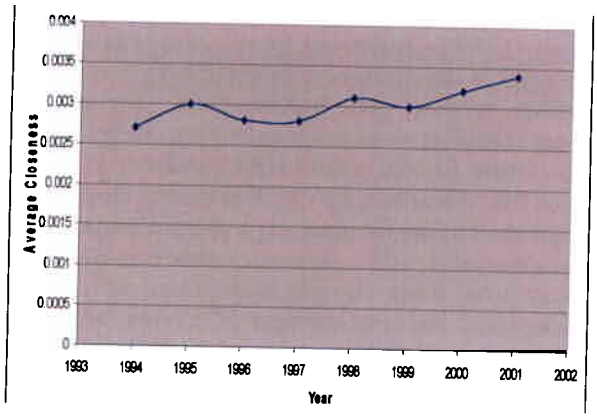

Figure 2. CUSUM Control Chart for Average Closeness of Al-Qaeda

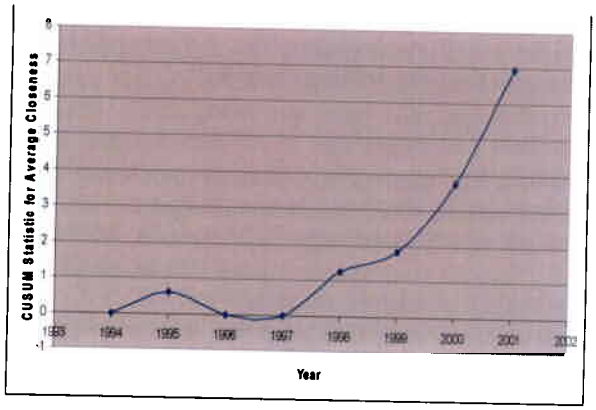

is used for increases and the other chart for decreases. Table 1 displays the CUSUM statistic values for detecting increases in a measure, while table 2 shows the values for decreases in a measure.

The control chart will signal a false alarm after 168 observations on average when the control limit is arbitrarily set to $h = 4$ (McCulloh, 2004). This corresponds to a probability of false alarm of

Table 1. CUSUM Statistic for Detecting Increases in a Network Measure

	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004
Number Agents	0.00	0.00	0.00	0.00	1.19	2.12	3.43	4.99	2.18	0.00	0.00
Average Centrality	0.00	0.60	0.00	0.00	1.19	3.74	6.10	9.24	4.39	0.00	0.00
Average Betweenness	0.00	0.43	0.00	0.00	2.02	3.24	5.26	8.46	5.31	0.00	0.00
Average Closeness	0.00	0.59	0.00	0.00	1.25	1.84	3.74	6.95	3.61	0.00	0.00
Average Density	0.00	0.60	0.00	0.00	1.19	3.74	6.10	9.24	4.39	0.00	0.00

Table 2. CUSUM Statistic for Detecting Decreases in a Network Measure

	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004
Number Agents	0.00	0.00	0.37	0.37	0.00	0.00	0.00	0.00	1.81	6.68	16.48
Average Centrality	0.34	0.00	0.00	0.15	0.00	0.00	0.00	0.00	3.85	12.58	27.15
Average Betweenness	0.00	0.00	0.56	0.72	0.00	0.00	0.00	0.00	2.15	9.07	19.16
Average Closeness	0.37	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.34	10.57	26.00
Average Density	0.34	0.00	0.00	0.15	0.00	0.00	0.00	0.00	3.85	12.58	27.15

ISBN 978-1-934808-00-9

less than 1%. It can be seen in Table 1 that the CUSUM statistic exceeds the control limit of 4 and signals that there might be a significant change in the Al-Qaeda network in either 2000 or 2001 for all five measures. Therefore, an analyst monitoring Al-Qaeda would be alerted to a critical, yet subtle change in the network prior to the September 11 terrorist attacks.

The CUSUM control chart also has a built in feature for determining the most likely time that the change occurred. This time is identified as the last point in time when the CUSUM statistic is equal to 0. For all measures, this point in time is 1997. To understand the cause of the change in the Al-Qaeda network, an analyst should look at events occurring in 1997.

Several very interesting events related to Al-Qaeda and Islamic extremism occurred in 1997. Six Islamic militants massacred 58 foreign tourists and at least four Egyptians in Luxor, Egypt. Coalition forces deployed to Egypt in 1997 for a bi-annual training exercise were repeatedly attacked by Islamic militants. The coalition suffered numerous casualties and shortened their deployment. In early 1998, Zawahiri and Bin Laden were publicly reunited, although based on press release timings; they must have been working throughout 1997 planning future terrorist operations. In February of 1998, an Arab newspaper introduced the "International Islamic Front for Combating Crusaders and Jews." This organization, established in 1997, was founded by Bin Laden, Zawahiri, leaders of the Egyptian Islamic Group, the Jamiat-ul-Ulema-e-Pakistan, and the Jihad Movement in Bangladesh, among others. The Front condemned the sins of American foreign policy and called on every Muslim to comply with God's order to kill the Americans and plunder their money. Six months later the US embassies in Tanzania and Kenya were bombed by Al-Qaeda. Essentially, 1997 was possibly the most critical year in uniting Islamic militants and organizing Al-Qaeda for offensive terrorist attacks against the United States.

Table 2 shows similar results for detecting decreases in network measures. The control chart for all measures signals a possible network change in 2003. The most likely time that this change occurred was 2001, which corresponds to the U.S. invasion of Afghanistan.

CONCLUSIONS

Control charts are a critical quality engineering tool that assists manufacturing firms maintain profitability. This Al-Qaeda example demonstrates that social network monitoring could enable analysts to detect important changes in terrorist networks. Furthermore, the most likely time that the change occurred can also be detected. This paper does not mean to imply that the CUSUM statistic is the answer to social network monitoring. The CUSUM is simply used to illustrate the usefulness of a statistical process control chart for monitoring social networks. There are many other control chart statistics that could be used to monitor a dynamic social network. More research in this area is needed to characterize the nature of dynamic social networks, and to identify what statistics are best to minimize the probability of false alarms and increase the speed of detecting changes.

Future research in social network monitoring should also focus on friendly organizations as opposed to terrorist networks. Although understanding terrorism is an important application of social network monitoring, there is always a great amount of unknown information in terrorist organizations. The terrorist are just not good at filling out the surveys required to understand the true social network. Friendly organizations, on the other hand, provide researchers the opportunity to monitor all social connections, and through surveys and interviews, understand why those connections exist and when they change.

With a better understanding of social network monitoring, there is a wide range of potential applications. Intelligence analysts can better monitor terrorist organizations. Military commanders can have improved situational awareness of their units by monitoring communications among subordinates and linking the communication network to morale and motivation. Civilian business leaders can monitor the success and progress of strategic combinations, such as mergers and acquisitions. Most application areas of dynamic network analysis can make use of social network monitoring to detect and identify changes in a dynamic social network.

ACKNOWLEDGEMENT

This work was supported in part by the Office of Naval Research (ONR), United States Navy Grant No. N00014-02-10973 on Dynamic Network Analysis, the Army Research Labs Grant No. DAAD19-01-2-0009 , the Air Force Office of Sponsored Research (MURI: Cultural Modeling of the Adversary Organization, 600322), and the NSF IGERT program in CASOS (DGE-9972762). Additional support was provided by CASOS and ISR at Carnegie Mellon University. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official

policies, either expressed or implied, of the Office of Naval Research, the Army Research Labs, the Air Force Office of Sponsored Research, the Department of Defense, the National Science Foundation, or the U.S. government.

REFERENCES

Carley, K. M. (2003). Dynamic network analysis. *Dynamic social network analysis: Workshop summary and papers*: 133-145. P. Pattison (Ed.), Washington D.C.: The National Academies Press.

Carley, Kathleen & Reminga, Jeffrey & Borgatti, Steve. (2003). Destabilizing Dynamic Networks Under Conditions of Uncertainty. *International Conference on Integration of Knowledge Intensive Multi-Agent Systems, 2003. IEEE KIMAS, Boston MA*, Boston MA: IEEE KIMAS.

Carley, Kathleen. (2003). Destabilizing Terrorist Networks. *Proceedings of the 8th International Command and Control Research and Technology Symposium. Conference held at the National Defense War College, Washington DC. Evidence Based Research, Track 3, Electronic Publication*, WebSite:http://www.dodccrp.org/events/2003/8th_ICCRTS/pdf/021.pdf

Carley, Kathleen & Lee, Ju-Sung & Krackhardt, David. (2001). Destabilizing Networks. *Connections*, 24(3), 31-34.

Frantz, Terrill & Carley, Kathleen. (2005). A Formal Characterization of Cellular Networks. *Carnegie Mellon University, School of Computer Science, Institute for Software Research International, Technical Report CMU-ISRI-05-109*

Ghosh, B.K. (1970). *Sequential Tests of Statistical Hypotheses*, Addison-Wesley Publishing Company, Inc., Reading, Massachusetts.

Govindarajulu, Z. (1981). *The Sequential Statistical Analysis*, American Sciences Press, Inc., Columbus Ohio.

Lorden, G. (1971). Procedures for Reacting to a Change in Distribution. *Annals of Mathematical Statistics* **42**, pp. 1897-1908.

Lucas, J.M. (1982). Combined Shewhart-CUSUM Quality Control Schemes. *Journal of Quality Technology* **14**, pp. 51-59.

Marquand, Robert (2001). The tenets of terror. *Christian Science Monitor*,18 Oct 2001.

McCulloh, I. (2004). *Generalized Cumulative Sum Control Charts.* Masters Thesis, Florida State University, Tallahassee, FL.

Montgomery, D.C. (1991). *Introduction to Statistical Quality Control*, 2nd Edition, John Wiley and Sons, New York.

Moustakides, G.V. (1986). Optimal Stopping Rules for Detecting Changes in Distributions. *Annals of Mathematical Statistics* **14**, 1379-1387.Moustakides, G.V. (1998). Quickest Detection of Abrupt Changes for a Class of Random Processes. *IEEE Transactions* **44**, pp. 1965-1968.

Moustakides, G.V. (2004). Optimality of the CUSUM Procedure in Continuous Time. *Annals of Statistics.* **32**, (to appear).

Neyman, J. and Pearson, E.S. (1933). On the Problem of the Most Efficient Tests of Statistical Hypotheses. *Philosophical Transactions Royal Society Series A.* **231**, pp. 289-337.

Page, E.S. (1954). Continuous Inspection Schemes. *Biometrika* **41**, pp. 100-115.

Page, E.S. (1961). Cumulative Sum Control Charts. *Technometrics* **3**, pp. 1-9.

Pignatiello, J.J., Jr. and Samuel, T.R. (2001). Estimation of the Change Point of a Normal Process Mean in SPC Applications. *Journal of Quality and Technology*. **33**, pp. 82-95.

Shiryayev, A.N. (1996). Minimax Optimality of the Method of Cumulative Sums (CUSUM) in the Case of Continuous Time. *Russian Mathematics Survey* **51**, pp. 750-751.

Siegmund, D. (1986). Boundary Crossing Probabilities and Statistical Applications. *Annals of Mathematical Statistics* **14**, pp. 361-404.

Topper, Curtis & Carley, Kathleen. (1999). A Structural Perspective on the Emergence of network Organizations. *Journal of Mathematical Sociology*, 24(1), 67-96.

Wald A. (1945). Sequential Tests of Statistical Hypotheses. *Annals of Mathematical Statistics* **16**, pp. 117-186.

Wald A. (1947). *Sequential Analysis*, Wiley, New York.

Wald, A. and Wolfowitz, J. (1948). Optimum Character of the Sequential Probability Ratio Test. *Annals of Mathematical Statistics* **19**, pp. 326-339.

Weiss, L. (1953). Testing One Simple Hypothesis Against Another. *Annals of Mathematical Statistics* **24**, pp. 273-281.

Weiss, L. (1962). On Sequential Tests Which Minimize the Maximum Expected Sample Size. *Journal of the American Statistical Association* **57**, pp. 551-557.

Vance, L.C. (1986). Average Run Lengths of Cumulative Sum Control Charts for Controlling Normal Means. *Journal of Quality Technology* **18**, pp. 189-193.

AUTHORS

Ian McCulloh is a Major in the U.S. Army. He's currently completing a PhD at Carnegie Mellon University.

Kathleen M. Carley is a professor of Computer Science at the Institute for Software Research, Carnegie Mellon University, where she directs the Center for Computational Analysis of Social and Organizational Systems (CASOS).

Matthew Webb is a 2nd Lieutenant in the U.S. Army. He co-authored this paper while a cadet at the U.S. Military Academy, Class of 2007.