Hubs, Authorities, and Networks: Predicting Conflict Using Events Data

Richard J. Stoll
Rice University
Department of Political Science
MS24
Box 1892
Houston, TX 77251-1892
Phone: 713-348-3362
FAX: 713-348-5273
stoll@rice.edu

Devika Subramanian
Rice University
Department of Computer Science
MS132
Box 1892
Houston, TX 77251-1892
Phone: 713-348-5661
FAX: 713-348-5930
devika@rice.edu

Paper prepared for presentation at the 2006 Annual Meeting of the International Studies Association, San Diego, CA, March 22-25, 2006. Portions of the work done in this paper were supported by the National Science Foundation under grant NSF-ITR 0219673.

Abstract

We conjecture that before states in a region engage in high levels of conflict, there will be noticeable changes in the pattern of events among the states in the region. If this conjecture is correct, then uncovering this pattern (or these patterns) could serve as an early warning indicator of the onset of serious conflict. We examine this conjecture using events data for the Levant region. We first use events data to create hubs and authorities scores for countries (note: hubs and authorities data are often used to identify key nodes on the Internet, with a hub being a website that links to a number of important websites, and an authority being a website that is linked to by a number of important websites). We use summary measures of the negative hubs and authorities scores and relate these changes to the outbreak of serious conflict in the region.

Introduction

Since September of 2001 terrorism has been a major focus of US foreign policy. This is appropriate. But the United States cannot afford to ignore the prospect of "ordinary" armed conflict between states. Such conflicts are not only dangerous and costly for the participants, but they may also threaten the interests of the United States. So it is important that we develop ways of predicting the outbreak of such conflicts. This paper is part of one attempt to do so. It is part of a larger effort to generate events data from online news sources and to use these data to explore a wide variety of questions in international relations.

Here we rely on events data collected by others (KEDS Project, 2006a) and concentrate on devising a means to predict the onset of serious conflict. We will use measures that are currently used to map the important nodes in the internet – hubs and authorities – to summarize the relationships within a region. We will then use these measures to see if they can alert us to the possible outbreak of serious conflict.

Early Warning of Conflict

We believe that the outbreak of a serious interstate conflict is not complete surprise. The specifics of when and how a conflict becomes militarized may be unanticipated, but not the fact that that there is conflict. Consider the Japanese attack on Pearl Harbor. There is no question that the *attack* was a surprise. But it was not a surprise that the United States and Japan were at odds with one another. The following item appeared in the *New York Times* in January, 1941:

> **Sees U.S.-Japanese Conflict**
>
> TOKYO, Tuesday Jan. 7 (AP)—The Japanese newspaper Kokumin predicted editorially today that "Japan and the United States will fight like wildcats" in the Pacific after collapse of Britain, adding that "it is safe to predict the Britain will be wiped off the European map by Autumn at the latest.
>
> "Even granting that the United States is not entering the European war, so long as she pursues her present pro-British activities she is in it," the editorial went on.
>
> "Had Germany been bellicose, she would have found more than enough reasons to declare war against the United States for past contemptuous utterances and illegal deeds. The same applies to Italy and Japan."
>
> (New York Times, 1941a).

The story does not contain a codeable event between Japan and the United States (see below). But it does illustrate that the Japanese newspaper considered it likely that Japan and the United States would become engaged in a conflict.

As well, on June 8, 1941 the *New York Times* reported that Rear Admiral Harry E. Yarnell who had just stepped down as commander of the US Asiatic Fleet, said in a speech at the Stevens Institute that the United States was facing its "third great crisis"[1] and that the United States should fight now (New York Times, 1941b). Again, this is not a codeable event, but it illustrates that there was recognition in both Japan and the United States that the states were likely to become involved in conflict against one another.

We believe that as states move towards serious conflict (disputes and wars) observers on both sides will sense what is happening and comment on it. But more importantly, in these situations the

[1] Yarnell thought the first two crises were the American Revolution and the Civil War.

states will take actions. And these actions will help us identify the conditions which signify that a serious conflict is imminent.

Events Data

Events data are intended to describe the lower-level interactions between actors (usually states). These data have the potential to allow researchers to track the low level interactions between states. They became popular during the 1970s and 1980s (for example see Azar, 1970 or McClelland, 1978). But with the ending of funding for the major events data collection projects,[2] interest in using these data began to die out both because of its seeming "obsolescence" (the ending years of these datasets grew more and more distant from the current year), and questions about the entire approach (McClelland, 1983).

In the 1990s events data collection was revived, but using a different approach (fore example. Schrodt, Davis and Weddle, 1994 or Bond, Jenkins, Taylor, and Schock, 1997). This revival was based on two developments. First, as the Internet developed, more and more news outlets created websites. These websites not only contained current news stories but in many cases archives of older stories.[3] Second, advances in computer science have facilitated the development of software to "read" news stories and extract events data. While this technology is by no means perfect, there are distinct advantages to using software instead of people to do events coding. For example, "[c]omputer programs do not get tired, bored, and distracted, and so in the long run the program would certainly outdo any human coder that would be feasible for a researcher to recruit" (King and Lowe, 2003, p.

---

[2] In that era, events data were generated by human coders reading news sources. This approach required large numbers of trained coders, and therefore required a great deal of funding.

[3] To give but one example, at the *New York Times* website (http://nytimes.com/) the archive extends back to the founding of the newspaper in 1851.

619). And it appears that computer programs perform as well as human coders (Schrodt and Gerner, 1994; King and Lowe, 2003). Software for extracting events data includes VRA Reader (Virtual Research Associates, 2006), and Tabari (KEDS Project, 2006).[4]

There are a number of potential problems with events data. For example, the news sources used may be incomplete or biased. As well, many important activities take place in secret and are not reported in news sources (McClelland, 1983). And while computer programs do not get tired (see above), they also do not possess the sophistication of the human mind. Consequently, these programs may fail to provide accurate coding in some circumstances. But nevertheless, the computer extraction of events data is a potentially powerful tool. It allows us – albeit imperfectly – to track and trace the actions of international actors.

There is an additional element to most event data schemes. Each scheme describes particular events that are to be coded. But in most cases, the scheme also has an associated scale that rates each event type in terms its degree of conflict or cooperation.[5] In this paper we will use events data coded according to the WEIS scheme. We will also use the conflict-cooperation scale developed by Goldstein (1992).

Hubs and Authorities[6]

Given individual events, how does one put them together to look for patterns and to predict the outbreak of serious conflict? There are two separate issues here. The first is to what extent individuals

[4] Part of our own work is to develop software to code events data; for information on our efforts to date see REDS, 2006.
[5] At the present time all scales for events data that are well-known reflect conflict and cooperation. But scales could be based on other dimensions.
[6] The discussion of hubs and authorities is taken from Subramanian and Stoll, 2004.

events should be aggregated across time. The second is how to characterize the relationship between states given the events in which they participate.

A cursory examination of the sources for events data makes it clear that foregoing aggregation and using single events is fraught with difficulties. Suppose we are studying the interaction of a pair of states. To supply a specific context, consider the United States and the Soviet Union during the Cuban Missile Crisis. What should we do when we code more than one event in a day? One problem is that many accounts do not contain sufficient temporal information to unambiguously order events during a day. But even if we could determine the specific time during the day that various events occurred, would this suffice? Probably not. Recall that during the crisis, within the US government there was concern that with the time difference between Washington and Moscow (8 hours) that it would be difficult for one of the governments to quickly respond to an individual event initiated by the other. So some degree of aggregation is usually required even if we have precise information on the timing of events. Various temporal aggregations have been used in studies with events data; periods of time range up to a year. But here we will use a much smaller unit of time. We will aggregate events into two week time periods.

After picking a temporal aggregation, there still remains the larger question of how to take the events data scores of a set of states and determine the relationships between them. We will use an approach that has become a popular way to evaluate the importance of websites. Data on the importance of websites are something that most of us use on a regular basis; this is part of what happens any time we use a search engine.

How would we recognize an important website? One way is to find websites that are referenced by a larger number of other important websites. Such websites are called *authorities*. These are websites that we would frequently visit. But there is another type of website that we would frequently visit. It is one that references a lot of important websites. That type of website is called a *hub*.

There is no single or best way to determine hubs and authorities on the Internet. If there was, there would only be a single search engine. But to illustrate the ways in which hubs and authorities might be identified, we follow the discussion in Chakrabarti et al. (1999). They suggest the following procedure to calculate hubs and authority scores for websites:

1. Look for the term of interest using a search engine and identify the first 200 webpages that are returned by the search.
2. Now include all the webpages that link to the first 200 webpages. Also include all the webpages that link from the first 200 webpages. This is called *the root set*.
3. Begin the process of calculating hub and authority scores. The initial hub score for a page is the number of pages in the root set that the page references. The initial authority score for a page is the number of pages in the root set that have links to that page.
4. Once all the initial scores are assigned, recalculate the hub and authority scores. Use the initial scores for each page as a weight for its contribution. For example, as noted in the initial assignment, the authority score of a page is equal to the number of pages that point to it. But in the next iteration the authority score of a page is equal to the sum of the hub scores of the pages that link to it.
5. Continue to iterate and update using the scores from the last iteration until the scores stabilize.[7]

This process of calculating scores is reasonable and straightforward. But in order to apply it to identifying linkages between states using events data, we introduce some additional elements to the

[7] Chakrabarti et al., 1999 say that stabilization occurs fairly quickly. They say that a root set of 3000 pages will stabilize after about five iterations [note: we cannot provide page references to this piece because we accessed it through the web].

discussion. We will use concepts from graph theory to provide a more formal context for our discussion of hubs and authorities. We draw from link analysis to analyze the relationships. And we also have to deal with the fact that relationships between states as measured with events data are more complicated that the simple presence or absence of a link as is the case with webpages.

Graph Theory and Link Analysis

Graph theory (Harary, Norman, and Cartwright, 1965) deals with the representation and study of objects (called nodes) and the connections between them (called edges). Edges can simply indicate there is a connection between nodes, or they may be directed (i.e., indicate that the link goes from one node to the other). A collection of nodes and edges is called a graph. If a pair of nodes can be connected by more than one edge, this is termed a multi-graph.

In this research we are dealing with directed multi-graphs in which states are nodes and the event streams connecting them are edges. We use separate directed edges to indicate the cooperative and conflictual event streams for one state to the other. So there can be 4 edges connecting two state nodes, A and B: cooperation from A to B, conflict from A to B, cooperation from B to A, and conflict from B to A.

Link analysis is used to analyze directed multi-graphs. It is utilized to identify important nodes and edges based on the structure of the graph. It has been used to analyze the structure of links between webpages (Kleinberg, 1998) and the structure of links between citations (Osareh, 1996).[8]

Calculating Hubs and Authorities Scores

[8] A good review paper that cites a number of important papers would be a significant hub. An important paper that is cited by many other papers (hubs) would be a significant authority.

To build our graphs we begin by calculating *Aw*, the authority weight of a state (indicating the events directed to that state) and *Hw* the hub weight of a state (indicating the events that state initiated against other states). We do separate hubs and authority calculations for the positive (cooperative) and negative (conflictual) events.

Between each pair of states there are a total of four hubs and authorities scores. In the calculations *e(q,p)* is the scale score associated with an event in which state *q* is the initiator and state *p* is the target (i.e., the degree to which the event is cooperative or conflictual). Positive authority weights are defined in terms of the positive hub weights, and the positive hub weights are defined in terms of the positive authority weights. A similar relationship holds between the negative authority and hub weights. The equations that define the four hubs and authorities scores are:

$$posAw(p) = \sum_{q:(q,p)\in E} e(q,p) * posHw(q)$$
$$posHw(q) = \sum_{p:(q,p)\in E} e(q,p) * posAw(p)$$
$$negAw(p) = \sum_{q:(q,p)\in E} e(q,p) * negHw(q)$$
$$negHw(q) = \sum_{p:(q,p)\in E} e(q,p) * negAw(p)$$

Hub and authority weights are computed by a fix-point iterative method. Note that by reformulating the above equations in matrix form, with *PosAw*, *PosHw* and *NegAw* and *NegHw* standing for the vector of hub and authority weights and *W* for the interaction matrix (edge weights in the interaction graph), we see that

$$PosAw = W \;\; * PosHw$$
$$PosHw = W^T * PosAw$$
$$NegAw = W \;\; * NegHw$$
$$NegHw = W^T * NegAw$$

which after algebraic simplification yields:

$$PosAw = (WW^T)PosAw$$
$$PosHw = (W^TW)PosHw$$
$$NegAw = (WW^T)NegAw$$
$$NegHw = (W^TW)NegHw$$

PosAw and *PosHw* are eigenvalues of WW^T and W^TW respectively; *NegAw* and *NegHw* are defined analogously. These quantities can be computed by standard eigenvector computation techniques (Golub and Van Loan, 1996) that are part of statistical packages such as Matlab and R.

Measures and Analysis: Hubs, Authorities and Conflict

Given a set of hubs and authorities scores we need a way to construct measures that summarize these scores to predict the onset of significant amounts of conflict. Because our interest is in the prediction of conflict for this work we will only use the negative hubs and authorities scores for our measures. In future work we will develop additional measures that will utilize both positive and negative hubs and authorities.

Our first measure is quite simple. It is the number of negative components[9] in a time period. Each connected component in an interaction graph is a set of nodes that are reachable from each other by directed paths. This measure taps the number of separate or independent units.

The second measure is an adaptation of a well-known measure of the distribution of capability, CON (Singer, Bremer, and Stuckey, 1972). We construct a measure based on the distribution of states across negative components. Consider an instance in which every state was in its own component. If that was the case, it seems implausible that there would be a significant amount of conflict among

[9] "Negative components" means the number of components calculated from the negative hubs and authorities scores.

those states; no two states are in a single component defined by their conflictual events. Now consider the opposite situation. Virtually all the states are in the same component as defined by the conflict events between them. If that were the case, then we would expect a great deal of conflict. Scoring each negative component by the number of states in it, we calculate the standard deviation of the component scores. To normalize this, we divide by the maximum possible standard deviation given the number of components. This would occur if all the states were in the same component.[10]

While our first measure might be viewed as a control variable, the second measure has clear substantive meaning. As more and more states cluster together base on their negative interactions, we would expect that there would be higher levels of conflict. And if our work is to have any value to serve as a way to alert others that conflict is going to erupt, the measures we develop should be useful in anticipating future conflict.

As for our dependent variable, we aggregate all WEIS events using the Goldstein scale (1992) of conflict and cooperation to create a single measure of the degree to which the time period was cooperative (a positive total) or conflictual (a negative score). Our time period runs from April 15, 1979 through December 31, 1997. These data are taken from the Levant data set produced by KEDS (2006a). This is a computer-generated events data set that contains all coded events involving the following entities: Egypt, Israel, Jordan, Lebanon, Palestinians, Syria, USA, and USSR/Russia.[11] As noted above, we aggregate these events into two-week time periods.

It is readily apparent that these aggregated events data are likely to exhibit the typical characteristics of time series data. The time series nature of the data need to be evaluated and, if necessary, the appropriate steps taken to insure valid inference. An examination of the autocorrelation

[10] Obviously this hypothetical situation could not arise in practice. If all states were in the same component, then there would be only one observation and the standard deviation would be zero. But nevertheless, conceptually this is a reasonable way to normalize the observed standard deviation.

[11] To be clear, the Levant dataset includes events that involve any one of these states, with any other state in the world; for example, the Levant dataset would include events between Egypt and Brazil.

and partial autocorrelations functions of the dependent variable revealed that is was very likely that an AR(1 2) process was present. Consequently, all analysis were done using ARIMA with the AR(1 2) process included in the analysis. Table 1 displays the summary statistics of the variables to be used in our analysis.

Table1
Descriptive Statistics of Variables Used in Analysis

Variable	**N**	**Mean**	**Std. Dev.**	**Minimum**	**Maximum**
Sum of Scaled Events scores	487	-129.826	157.186	-1189.9	529.7
Number of Components	487	17.099	4.599	5	34
Relative Std. Dev. Of Component Scores	487	.251	.119	0	.600

Notice that the mean value for the sum of events scores is negative. We believe that this is due to the fact that the states of the Levant region were often engaged in conflict. For example, using Bennett's (1998) list of enduring rivalries, all seven of the states that KEDS codes in the Levant data participate in rivalries during at least a portion of the time period under study. The Palestinians are the eighth actor in the Levant dataset; since they were not considered a state, they cannot be part of an enduring rivalry. However, it is not much of a stretch to believe that they too were engaged in a significant amount of conflict during that time.[12] This makes this group of states important to study. But it also should serve as a caution; what we learn from studying the interactions of a particularly conflict-prone set of states may not apply to other groups of states.

[12] One suspects that the researchers at KEDS chose to collect data on these eight actors precisely because they were experiencing high levels of conflict.

Figure 1 plots the summed scale events scores through time. As can be seen, most of the bi-weekly scores lie between 500 and -500 with a handful taking on more extreme values. Note also that most of the scores lie below zero. This reinforces what we saw from the descriptive statistics (and the discussion of enduring rivalries); this is a group of states that engages in a lot of conflict.

Figure 1

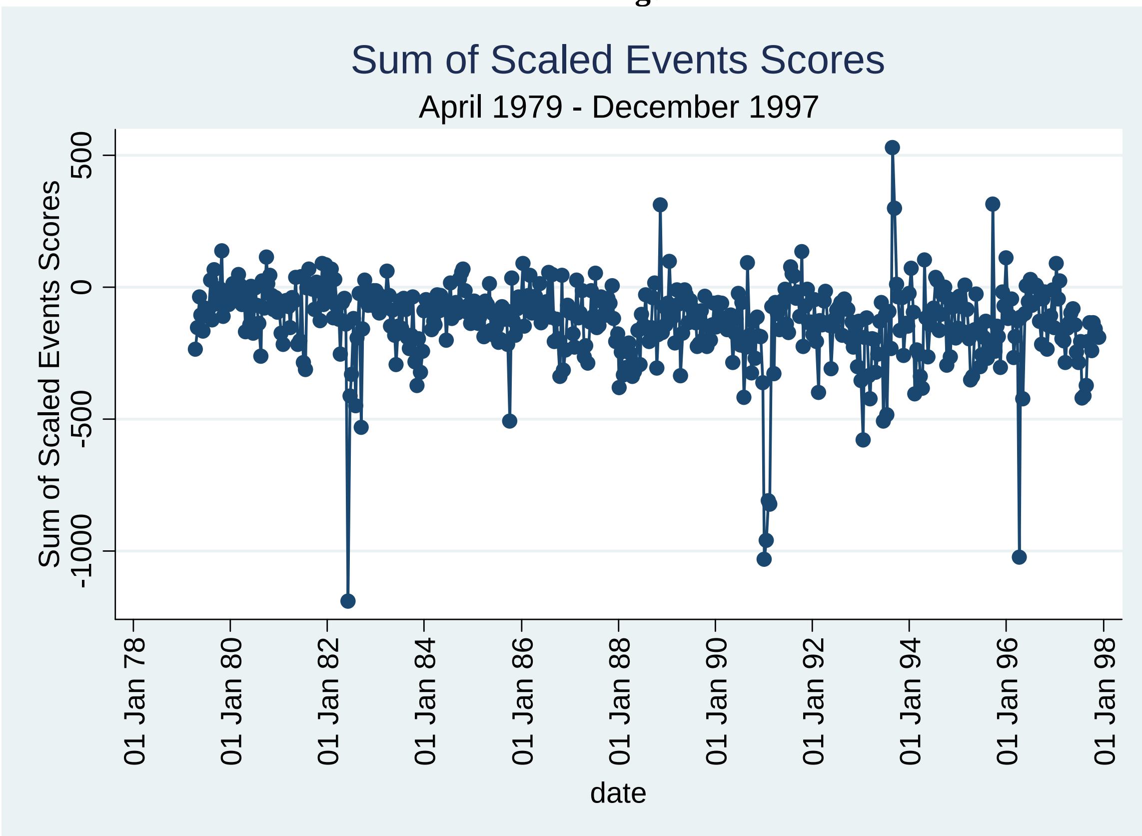

Figure 2 displays the plot of the number of components through time. Visually the pattern of values for the number of components appears to have fewer extreme spikes than the summed event scores shown in Figure 1. Almost all of the observations lie between 10 and 30.

Figure 2

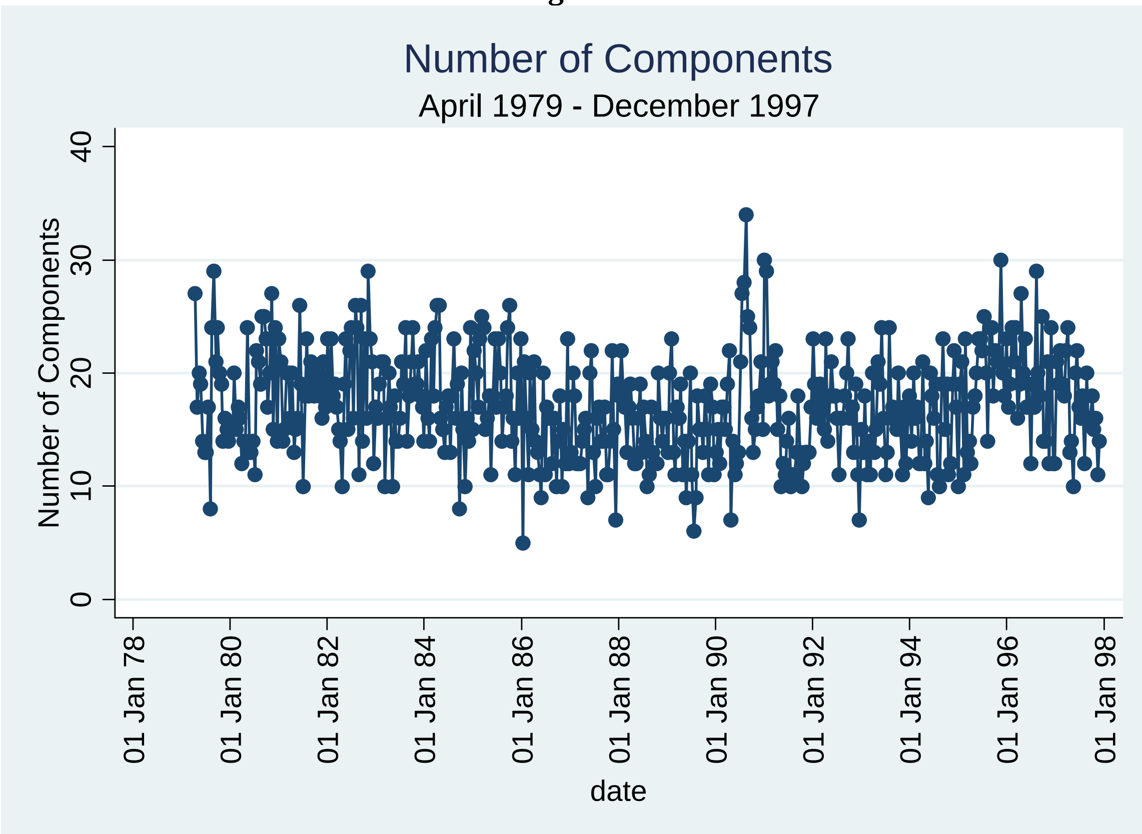

Finally we plot the relative standard deviation of component scores in Figure 3. Unlike the number of components, there appears to be a bit of a relationship between these scores and time. This makes sense if one considers how this variable is created. Components are scored by the number of states they contain. This in turn is a function of the amount of conflict between these states. We would expect that conflict would persist through time. This in turn would lead to the members of the components stay the same through time and consequently the standard deviation at one point in time would be related to the standard deviation in earlier points in time.

Figure 3

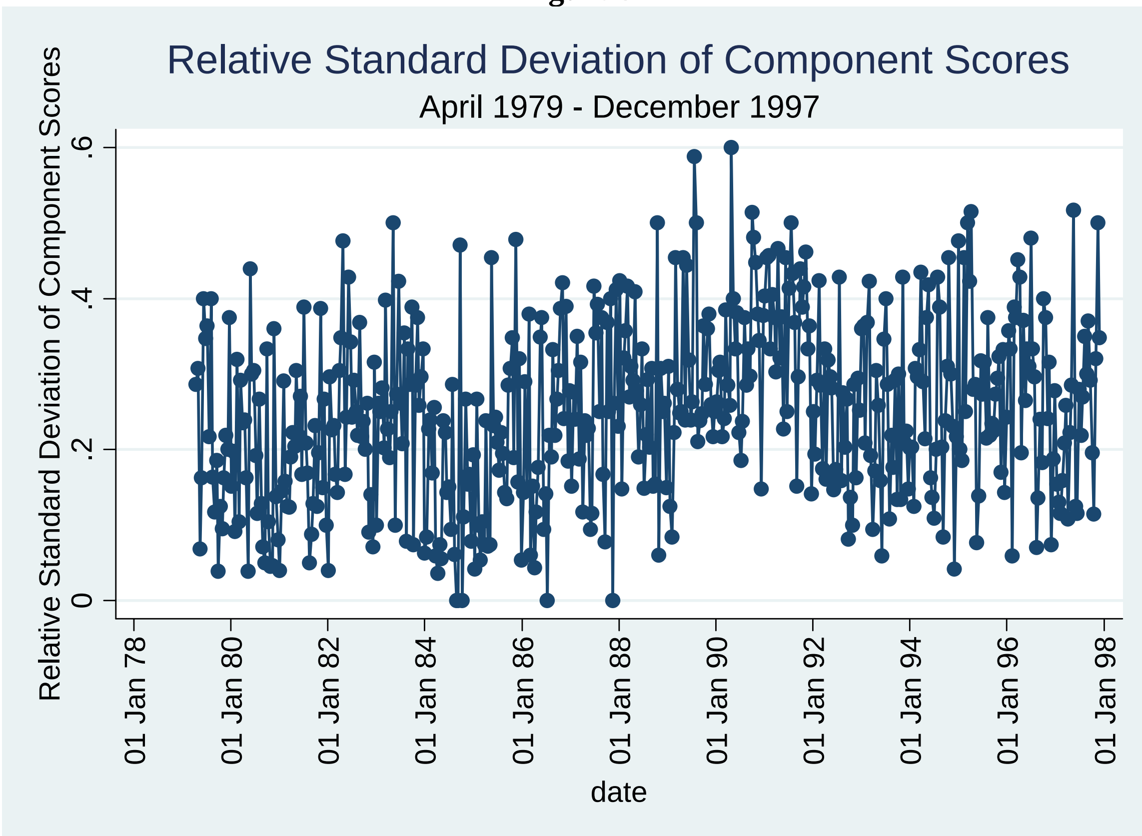

Predicting Events in the Levant

We now turn to the analysis predicting the summed scaled events scores. As noted, we analyze these data using ARIMA with two autoregressive lags. Table 2 displays the results of the analysis. We lagged the predictor variables for several reasons. First, unless we lag the predictors we cannot be sure that any relationship we observe runs from the predictor variable to the dependent variable; if all variables are measured at the same time, it is possible that the relationship runs in the reverse direction. Second, ultimately we want to be in a position to provide early warning of the onset of conflict. In order to do this the variables we use to provide this prediction must show the signs of impending conflict before it occurs. To be honest, we experimented a bit with different lags and came to the conclusion that the best solution was to use a lag of 3 bi-weeks for our independent variables.

Table 2
ARIMA Analysis Predicting Sum of Scaled Event Scores

```
Sample:  4 to 487                               Number of obs      =       484
                                                Wald chi2(4)       =    208.93
Log likelihood = -3088.697                      Prob > chi2        =    0.0000
```

	Coefficient	Std. Err.	z	P>|z|	[95% Conf.	Interval]
Number of Components (lag of 3)	-4.457614	1.602508	-2.78	0.005	-7.598472	-1.316756
Components: Relative S.D. (lag of 3)	-189.9823	69.36708	-2.74	0.006	-325.9393	-54.02532
Constant	-6.056485	42.14927	-0.14	0.886	-88.66753	76.55456
AR(1)	.2826723	.0319557	8.85	0.000	.2200402	.3453043
AR(2)	.1695744	.043734	3.88	0.000	.0838573	.2552915
/sigma	142.9471	2.223844	64.28	0.000	138.5884	147.3057

Note: Analysis conducted using Stata/SE 9.1

The results indicate that there is some relationship between our measures based on the negative hubs and authorities scores and the amount of conflict. Keeping in mind that conflict events have negative scale scores, we see that high numbers of components are associated with high levels of future conflict. Large values for the relative standard deviation of the components are also associated with high levels of future conflict.

As noted above in the discussion of the construction of these two variables, we view the number of components as a bit of a control variable. So we will focus our discussion on the interpretation of the relative standard deviation. If the relative standard deviation is large, this indicates that a small number of the components contain a large number of actors. The ARIMA results indicate large amounts of conflictual events are preceded by a clustering of actors into mutually reinforcing negative interactions. We think this is a potentially interesting finding.

As an additional measure to check on our results, we examined the residuals from the analysis. We examined those observations for which the absolute value of the residual was 400 or greater. To be sure, there is some element of arbitrariness in this threshold, but the value is about 2.5 times the size of the standard deviation of the dependent variable. Table 3 displays the date (i.e., the first day of the two-week time period), the residual, and one or more events involving the states in the Levant that we feel account for the larger than expected positive or negative score for that time period.[13]

Table 3
Examination of Large Residuals from ARIMA Analysis

Date	Observed Value	Predicted Value	Event
06 Jun 82	-1189.9	-127.8	Israeli invasion of Lebanon.
13 Nov 88	313.6	-141.2	Palestinian National Council endorses UN Resolution 242
06 Jan 91	-1030.8	-215.4	Gulf War.
20 Jan 91	-960.9	-411.4	Gulf War.
03 Mar 91	-73.8	-505.2	Shiite, Kurdish insurrection in Iraq. General strike in Israeli occupied territories.
17 Jan 93	-579.3	-178.0	US-Iraq conflict in no-fly zone. Hizbollah-Israel conflict in Lebanon.
29 Aug 93	529.7	-138.3	Israel & PLO sign agreement on limited self-rule in Gaza, Jericho.
24 Sep 95	316.9	-164.8	Israel & PLO sign Taba agreement on self-rule in West Bank
07 Apr 96	-1023.2	-179.2	Israel launches Operation Grapes of Wrath in Lebanon (Apr 11 – Apr 27).

Table 3 is presented primarily to allow the reader to make a better assessment of the fit of the model, but we will make a few comments. Most of the large residuals occur on weeks in which there was more conflict than cooperation. This is not too surprising; of the 487 cases, only 60 (a little over 12 percent) have positive values for the sum of the scale scores. Six of the cases involve Israel, and

[13] The events were identified by consulting various years of the International Institute for Strategic Studies publication *Strategic Survey*.

one can argue that some of the key events involving Israel were unanticipated, for example because the parties choose to keep the actions that led up to the key event a secret (for example, the Israelis did not announce Operation Grapes of Wrath until it started). While this may be true, if this leads to a large number of mis-predictions it does not bode well for our overall efforts. If most significant negative (or positive) time periods are not preceded by events that "point towards them," we will not be able to anticipate and predict them.

Discussion

Some might argue that the relative standard deviation measure is simply a re-labeling of the sum of the events (i.e., large standard deviations based on negative hubs and authorities scores are really the same thing as a large negative sum of events scores). While this argument is worth thinking about, we are not persuaded that this is all that we are doing. First of all, empirically the correlation between the sum of the scale scores and the relative standard deviation is not that great (the correlation is about -.31). If the two variables were measuring the same thing, we would expect to find a higher correlation. Second, conceptually the two measures have different foundations. The sum of event scale scores is based on the events that occur. Although events are part of the relative standard deviation measure, the measure itself is really about the grouping of states.

The argument does raise an interesting question. Look at the groupings of states that make up a particular value of relative standard deviation measure at time t-3, is it the case that later these states are the ones that engage in conflict events at time t? That is a question that we cannot answer at the moment but we will explore in the future.

Summary

We are in the midst of a revival of work using events data. Online news sources and advances in computer programming will make it possible for individual researchers to generate their own events data. While we are not quite there yet, that day will come and come soon. But what should we do with this new data source?

This paper is one attempt to answer that question. We attempt to predict the behavior (particularly conflict) of the countries in the KEDS Levant dataset. We do this by using a concept that is in current use as a way to study the Internet: hubs and authorities. These are ways to measure the importance and centrality of websites. We use these concepts as well as some ideas from graph theory to build measures of the degree to which states form components based on their negative behavior to one another. And we use these measures to predict the aggregate behavior of the actors in the Levant.

Although our success in this paper is modest, we believe that the overall approach does hold some promise. We believe that the interactions of states (as traced through events data) can be used to anticipate the outbreak of serious conflict. But it will be necessary to be creative in both creating and analyzing these data. This paper is a first step on this path.

References

Azar, Edward. 1970. Analysis of International Events. *Peace Research Reviews* 4,1: 1-113.

Bennett, Scott. 1998. Integrating and Testing Models of Rivalry Duration. *American Journal of Political Science* 42,4: 1200-1232.

Bond, Doug, J. Craig Jenkins, Charles L. Taylor and Kurt Schock. 1997. Mapping Mass Political Conflict and Civil Society: The Automated Development of Event Data. *Journal of Conflict Resolution* 41,4 :553-579.

Chakrabarti, Soumen, Byron Dom, David Gibson, Jon Kleinberg, S. Ravi Kumar, Prabhakar Raghavan, Sridhar Rajagopalan and Andrew Tomkins. 1999. Hypersearching the Web. *Scientific American*, June 1999. http://www.cs.cornell.edu/home/kleinber/sciam99.html. Accessed 2/15/2006.

Goldstein, Joshua. 1992. A Conflict-Cooperation Scale for WEIS Events Data. *Journal of Conflict Resolution* 36,2: 369-385.

Golub, Gene and Charles Van Loan. 1996. *Matrix Computations*. Baltimore, MD: Johns Hopkins.

Harary, Frank, Robert Norman, and Dorwin Cartwright. 1965. *Structural Models: An Introduction to the Theory of Directed Graphs*. New York: Wiley.

International Institute for Strategic Studies. Various Years. *Strategic Survey*. London, UK: Oxford University Press.

KEDs Project. 2006a. Levant Data Set. http://www.ku.edu/~keds/data.dir/levant.html. Accessed 3/16/2006.

KEDS Project. 2006b. TABARI: Text Analysis By Augmented Replacement Instructions. http://www.ku.edu/~keds/software.dir/tabari.info.html. Accessed 2/15/2006.

King, Gary and Will Lowe. 2003. An Automated Information Extraction Tool For International Conflict Data with Performance as Good as Human Coders: A Rare Events Evaluation Design. *International Organization*, 57, 3: 617-642

Kleinberg, Jon. 1999. Authoritative Sources in a Hyperlinked Environment. *Journal of the Association for Computing*. 46, 5: 604-632.

McCleary, Richard and Richard Hay. 1980. *Applied Time Series Analysis for the Social Sciences*. Beverly Hills, CA: Sage.

McClelland, Charles. 1978. World Event/Interaction Survey, 1966-1978. Ann Arbor, MI: Inter-university Consortium for Political and Social Research.

McClelland, Charles. 1983. Let the User Beware. *International Studies Quarterly* 27, 2: 169-177.

New York Times on the Web. 2006. http://nytimes.com/. Accessed 2/14/2006.

New York Times. 1941a. Sees U.S.-Japanese Conflict. *New York Times*. January 7, 1941, p. 3.

New York Times. 1941b. Yarnell Asks U.S. to Enter War Now. *New York Times* June 8, 1941, p. 28.

Osareh, Farideh. 1996. Bibliometrics, Citation Analysis and Co-Citation Analysis: A Review of Literature I. *Libri* 46: 149-158

Rice University Events Data System. 2006. http://hyperion.cs.rice.edu/general/default.htm. Accessed 2/15/2006.

Schrodt, Philip A., Shannon G. Davis and Judith L. Weddle, 1994. Political Science: KEDS—A Program for the Machine Coding of Event Data. *Social Science Computer Review* 12(3): 561-588.

Schrodt, Philip and Deborah Gerner. 1994. Validity assessment of a machine-coded event data set for the Middle East, 1982-92. *American Journal of Political Science* 38,3: 825-854.

Singer, J. David, Stuart Bremer, and John Stuckey. 1972. Capability Distribution, Uncertainty, and Major Power War, 1820-1965. In Bruce Russett (ed.) *Peace, War, and Numbers*. Beverly Hills, CA: Sage. pp. 19-48.

Subramanian, Devika and Richard J. Stoll. 2004. Four Way Street? Saudi Arabia's Behavior Among the Superpowers, 1966-1999. A study prepared for The Energy Dimension in Russian Global Strategy. Petroleum Energy Center of Japan and the James A. Baker III Institute for Public Policy.

Virtual Research Associates. 2006. Virtual Research Associates. http://vranet.com/. Accessed 2/15/2006.