Events, Patterns, and Analysis: Forecasting International Conflict in the Twenty-First Century

Devika Subramanian
Department of Computer Science
devika@rice.edu

Richard J. Stoll
Department of Political Science
stoll@rice.edu

Note: both authors are also affiliated with the Center for the Scientific Study of International Conflict and Cooperation

Rice University
Box 1892
Houston, TX 77251-1892

The research discussed in this paper is funded by the National Science Foundation, and INTEL.

This position paper was prepared for the workshop "Computer-Aided Methods for International Conflict Resolution" held at the Austrian Research Institute for Artificial Intelligence, Vienna, Austria, October 25-26, 2002.

This workshop is sponsored by the Austrian Federal Ministry for Education, Science, and Culture.

1. Project Description

Our research takes advantage of the vast amount of information that exists on the world wide web. In particular, we will make use of information from sites maintained by media outlets. We will use automated event extraction programs (Gerner et al., 1994) to construct events data, and then use the resulting time series of events to generate predictions about the onset and progression of serious international conflict. We hope to be able to predict the onset of these conflicts four to eight weeks in advance. Specifically, the goals of our research are:

- To design information extraction techniques and build events data sets for use by the entire scientific community.
- To use these events data and develop the algorithmic base for making predictions about the onset of serious conflict.
- To construct explanatory models, building on the existing findings from the scientific study of international relations (Vasquez, 2000).

Figure 1 provides a visual display of how we intend to organize and index the events for both retrieval and analysis.

Figure 1. Representation of the Data Gathering and Analysis Process

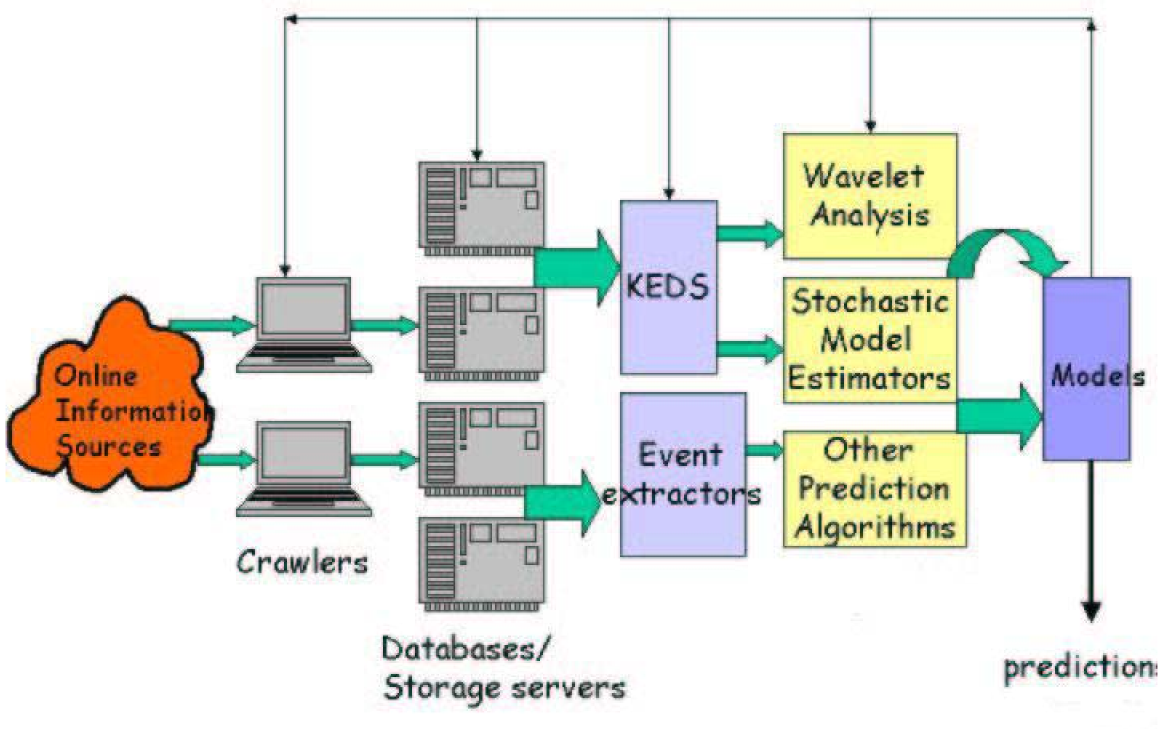

Timely warning of the outbreak of serious conflict can be a key element in conflict resolution. Early warning can provide the time for state and non-state actors to intervene and prevent the outbreak. Thus, we feel our work can be of potential value to the conflict resolution process, even though the focus of our research is the outbreak and evolution of conflict.

2. Components of the Research Program

2.1 Creating Events Data

In the social sciences, critical information for research rarely appears in a useful form. Instead, social scientists must engage in the process of *data making* (Singer, 1965). Events data -- nominal or ordinal codes recording the interactions between international actors -- are a prime example of this. A single record or case of events data consists of the actor (the entity that initiated the action), the action itself, the target of the action, and the date of the action. Most collections also associate a scale score with each event; this score represents the degree to which the event is cooperative or conflictual.

Until recently, the collection of these data was only possible through the use of teams of trained human coders who read through media sources to extract the appropriate information. This approach was both slow and expensive. Consequently the most widely used events data collections, COPDAB (Azar, 1980) and WEIS (Tomlinson, 1993), both cover only a limited period of time in the post-World War II era. We propose to automate the extraction process. This process will require a variety of state of the art tools. Some of these are already developed. One important tool is the set of software programs developed by the Kansas Events Data Project (KEDS -- see Gerner et al., 1994).

The creation of events data is basically a process of content analysis and involves three steps:

1. A source or sources of news about political interactions is identified. This could be an internationally-oriented newspaper such as The New York Times, a set of regional newspapers and newsmagazines, a news summary such as Facts on File or Deadline Data on World Affairs, or a newswire service such as Reuters or the Associated Press. Although they have not been used previously, and it poses additional complications, non-English language media sources can be used as well.
2. A coding system is developed, or a researcher may decide to use an existing coding system such as the World Events Interaction Survey (WEIS) the Conflict and Peace Data Bank (COPDAB) or the recently developed CAMEO (Conflict and Mediation Event Observations) coding scheme (Gerner et al., 2002). The coding system specifies what types of political interactions constitute an "event,"

identifies the political actors that will be coded (eg., whether nonstate actors such as international organizations and terrorist groups will be included in the data set), specifies the categories of events and their codes, and specifies any information to be coded in addition to the basic event. For example, the COPDAB data set codes a general "issue area"-- whether an action is primarily military, economic, diplomatic or one of five other types of relationship. WEIS, in contrast, codes for specific "issue arenas" such as the Vietnam War, Arab- Israeli conflict, and SALT negotiations.

3. In a machine-coding project, coding rules are implemented in a computer program such as KEDS or WebKB (Craven et al., 2000) by using extensive dictionaries (corpora) which identify actors and events. Corpus-driven techniques have proven to be much more robust than linguistically sophisticated methods that attempt to parse and understand natural language sentences. These techniques exploit implicit regularities in the structure of news stories. In fact, coding rules can now be learned directly from samples of suitably marked up documents (Craven et al., 2000), relieving us of the burden of formulating such extraction rules manually. We will use the tools made available by WebKB (Craven et al., 2000) to rapidly construct special purpose information extractors tuned to specific information sources.

KEDS has already been used to produce an extensive set of events data on several regions of the world (Schrodt and Gerner 1994, 1997), and it will be an important part of our effort to convert contemporary information into data. But we will also need to gather information from sources other than those that have been utilized previously by KEDS. We will extend the scope of KEDS-like data gatherers to more text sources on the World-Wide Web.

We also hope to develop the tools necessary to extract events data from non-English language sources. We do not underestimate the degree of difficulty of this task. However, the fact that we will be dealing with news media and focusing on event information will allow us to construct domain specific cross-lingual concept hierarchies (Yang et al. 1997) and use traditional copora-based information extraction techniques on these sources. We also anticipate collecting events data from multiple English language sources. A composite dataset will provide the most complete set of events, and also allow us to systematically evaluate the strengths and weaknesses of each source.

Event data sets can become quite large. We estimate that a global data set spanning the time period of the Cold War could encompass some 200 million events.

2.2 Event Data Patterns and Early Warning

Given the collection of events data, we will move to predict the outbreak of serious conflict. One particular form of serious international conflict that we seek to predict is militarized interstate disputes (Jones, Bremer, and Singer, 1996). Disputes are “united historical cases in which the threat, display or use of military force short of war by one member state is explicitly directed towards the government, official representatives, official forces, property, or territory of another state" (p. 268). These represent the most widely used set of conflict cases that fall short of war. We will also seek to predict aggregated measures of conflict from the events data themselves. That is, we seek to predict when the level of conflict reaches extremely high levels (although these levels may not reach the level of war). We think a reasonable goal is to identify outbreaks of serious conflict from four to eight weeks prior to onset. In our preliminary work, we have use the KEDS Gulf dataset, which runs from April 15, 1979 through March 31, 1999.[1]

The task of predicting the onset of a serious conflict such as a MID from a time series of scaled event scores is not easy. Conventional statistical and machine learning tools do not perform well on these data because they display significant levels of non-stationarity. Consequently, we believe that successfully dealing with these data requires several things. First, we need to select a period of aggregation for the data. Currently, we are using a two-week aggregation of all events, creating a single measure that summarizes the amount of conflict and/or cooperation that has taken place among the countries. This aggregation serves to average out a significant amount of noise in the data. Second, we need to use a different approach to analyzing these data. We are using multi-resolution wavelet analysis, both to demonstrate the degree of non-stationarity of the data, and then to uncover basic patterns in the time series. Even with the two-week aggregation, patterns in the overall data are not apparent from visual inspection. But the singularities (and their association with the outbreak of serious conflict) become much clearer after the data are subjected to wavelet analysis, as can be seen in Figure 2.

[1]This dataset is coded according to the WEIS coding scheme. It contains events involving the following countries: Saudi Arabia, Iran, Iraq, Kuwait, Oman, Yemen, Qatar, Bahrain, and the United Arab Emirates (KEDS, 2002).

Figure 2. Wavelet Analysis of Biweekly Gulf Data, 15 April 1979 - 31 March 1999

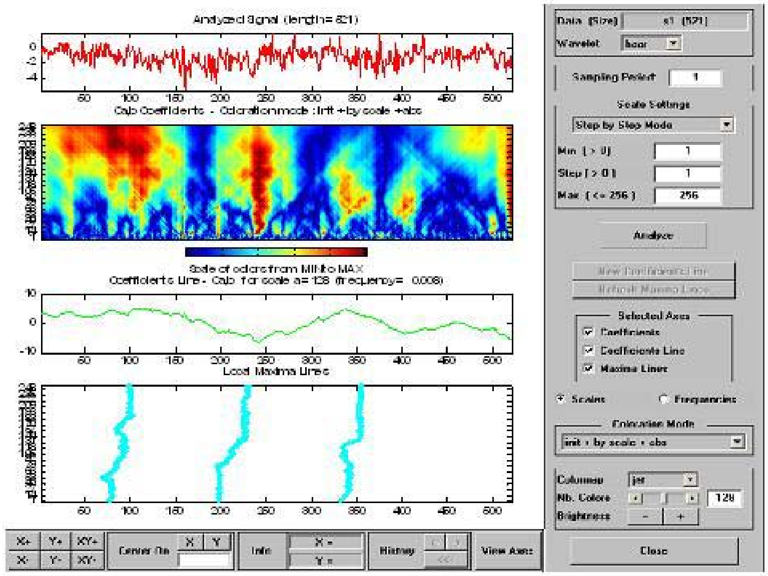

Note: Data are from the Kansas Events Data System (KEDS). Events involve the following countries: Saudi Arabia, Iran, Iraq, Kuwait, Oman, Yemen, Qatar, Bahrain, and the United Arab Emirates.

Our task is just beginning. Although we are encouraged with what we have uncovered so far, our analysis needs to be repeated and extended with larger datasets that span both a longer period of time, and a greater number of countries. This will allow us to determine if the same patterns exist in all regions, or whether we will have to devise different models for different sets of countries.

2.3 Modeling the Conflict Process

While uncovering patterns in the events data is an important step, it is not sufficient. We need to develop a deeper comprehension of the causes of conflict. To achieve this, we will develop models that will allow us to understand the forces that drive conflict to escalate or to de-escalate. A great deal of work has gone into predicting international conflict (although most of it has involved predicting the onset of war). In recent years, there has been a convergence of findings; good summaries of the state of the discipline are available in Bremer and Cusack, 1996, and Vasquez, 2000.

The factors that are most often associated with the escalation or de-escalation of conflict - usually in situations involving a pair of countries - are (Vasquez, 2000: 367):

- The presence of a territorial issue at stake between the two parties. Situations in which territorial issues are present are more likely to escalate; contiguous states are more likely to have territorial issues than states that do not share a common border.
- Some alliances promote war, while others promote peace. Alliances that settle territorial questions or which do not pose a great threat to a third party promote peace. Alliances that are formed by states that have fought a war recently and are dissatisfied with the status quo are likely to promote war (although the war does not break out immediately).
- Disputes involving states that are in an ongoing arms competition are more likely to escalate to war if nuclear weapons are absent.
- Disputes between parties that have had repeated conflicts are more likely to escalate to war.
- A dispute in which the parties bargain aggressively is likely to end in war. Across a series of disputes, if the parties use more and more aggressive bargaining strategies, a war becomes increasingly likely.
- Strong states are more war-prone than weak states. If the power balance between a pair of states rapidly moves towards parity, there is an increased chance of war.
- When major states establish norms or common expectations that reduce their freedom to take unilateral actions, the chances of war are reduced.
- Democratic states are very unlikely to go to war with one another.

These are all important findings that have turned up repeatedly in quantitative studies. They represent a set of empirical building blocks to increase our understanding. But we still need to formulate coherent models that combine sets of factors in a meaningful fashion.

While these findings all involve external factors, it is clear that foreign policy decisions are influenced by domestic factors as well. A good place to begin to integrate internal and external factors into a single model is the concept of the two-level game (Putnam, 1988). In terms of our work, we would integrate these two "games" using the following assumptions:

- In their external relations, state leaders are faced with a world resembling that described by classic realism. The primary factors that enter into a leader's calculations (in terms of using military force) are those mentioned in the summary of findings discussed earlier.
- At the same time, leaders seek to maintain their domestic political power, and these leaders must take this into account when they make all their decisions -- including foreign policy decisions.
- Domestic political groups have varying degrees of influence over the ability of leaders to maintain themselves in power.
- The interrelationship between domestic groups and leaders will differ depending on the domestic political system.

There are a variety of implications of these simple premises. First, leaders may be unable to pursue the best choice for foreign behavior because important domestic groups may not support this choice. Second, leaders may take actions primarily to build or maintain domestic support, but these actions may have significant foreign policy consequences. Third, if the leaders experience widespread domestic conflict, their responses to it will have consequences for the foreign policy of the state. We will construct an explicit model based on these implications, factoring in the nature of specific countries in the when necessary, but striving to keep the model as general as possible. Like the previous model, these factors are also static and do not change during the course of a crisis.

As we move forward with our exploration of these models, we will also move from variance theories to process theories (Mohr, 1982; Bremer, 1996). A variance theory deals with variables and efficient causes. Given two variables, X and Y, where Y is the variable to be explained, and X is a variable that has an impact on Y, the precursor variable, X, is idealized as being a necessary and sufficient condition for Y. As well, in a variance theory the time ordering among the precursor variables is assumed to be immaterial to the outcome. The model discussed above that uses state and dyadic characteristics fits comfortably within the variance theory perspective.

In a process theory, the focus is not on the variables, but on the linkages between them. A process theory deals with discrete states and events. A precursor variable, X, is a necessary condition for the outcome. Time ordering among precursor variables in a process theory is generally assumed to be critical in determining the outcome. The model mentioned earlier that is centered on the interaction of dispute strategies is an example of a model from the process theory perspective.

Perhaps a simple example will serve to highlight the difference between the two perspectives. Imagine the following ingredients: eggs, cheese, green peppers, onions, chives, and milk. If we try to predict the outcome (what item is being prepared) from the number and quantity of ingredients, we are using a variance theory. If we try to predict the outcome from the recipe, we are using a process theory. We believe that ultimately, we learn more about conflict by treating it as a dynamic situation to be modeled from the process perspective. To return to our simple example, it is only through the study of process -- the recipe -- that we can resolve whether the outcome of the list of ingredients noted above is scrambled eggs or an omelet.

In addition to building and testing models using well-established statistical analysis, we intend to apply approaches developed within the field of artificial intelligence. One promising avenue involves treating the evolution of conflict as a stochastic process to be modeled by a dynamic belief network (DBN; see Pearl, 1988). A belief network is a representation of a probability distribution on a set of variables from the problem domain at hand. It contains two parts: a qualitative representation and an associated quantitative representation. The qualitative part takes the form of an acyclic digraph. Each vertex in this graph represents a variable that can take one of a set of values. The set of arcs of the digraph defines a dependence relationship between the represented variables. Informally, the presence of an arc between two variables denotes a direct "influential" relationship between the linked variables, and the absence of an arc denotes that there is no direct influence between the two variables. The quantitative representation supplies a set of conditional probabilities for the links between variables. As discussed in Pearl (1988), there are several algorithms for probabilistic inference that can use this representation. We intend to evaluate the utility of these algorithms to help us understand the evolution of these conflicts.

2.4 Anticipated Problems

Our funded research does not begin until January, 2003. At that time, we will have a better understanding of the various pitfalls and roadblocks in our research. But in our preliminary efforts, we have developed a sense of where some of the problems in our research program will occur. Some of the issues that we will have to resolve include:

- *The identification of duplicate incidents*. There are great advantages to using multiple sources of information, but this increases the chances that we will uncover duplicate accounts of the same event. We have to develop a reliable way to screen out duplicate events.
- *Specific methods for wavelet analysis*. There are a variety of wavelets that can be used to analyze data. We need to develop a better sense of what patterns are likely to be uncovered by various wavelets. We need to have high confidence that our inferences are not driven by the use of a particular wavelet.
- *Prediction is different from post-diction*. Determining the presence of singularities with twenty years worth of data is different than predicting the onset of a singularity prior to its occurrence. A simple thought experiment illustrates the difference. Imagine you are walking; sometimes you go uphill, sometimes downhill. As you walk, the ground begins to slope downward. Is this a dramatic drop? The answer to this question becomes clear if you continue to walk; your judgment is based both on how far you go down, and a comparison of this situation with previous situations you have experienced. This is post-diction. But if you have to decide if there will be a great drop at the time you begin to go downhill, this prediction and a more difficult task.

3. The End Product: Real Time Early Warning and Understanding of Conflict

We have ambitious goals, but we believe they are feasible. First, we propose to build an extensive data archive of events between countries. Given the technologies available today, we believe that once we have the appropriate software up and running it will be possible to collect these dataset on a daily basis. By itself this will be a valuable resource for scholars, students, and policy makers. We will of course make these data available on the widest possible basis. Second, by examining patterns in events, we will develop early warning indicators of the onset of serious conflict. We believe that this will also be a valuable asset to these same three groups. It would be naïve to assume that we can have an immediate impact on the real world, but taking advantage of the world wide web, we believe that if we establish a track record of successful predictions, that impact may occur in the future. Finally, we will use the events data we generate as part of a major effort to model the process of conflict escalation. This will increase the knowledge base of the causes of conflict.

As we noted in the introduction, we believe that an early warning system can be an aid in conflict resolution. If we have high confidence that a conflict is going to break out, we can act to prevent it. We cannot be certain that such efforts will be successful, and conflict avoided, but we believe that a quality early warning system can make a significant contribution to successful conflict resolution.

References

Azar, Edward. 1980. The Conflict and Peace Data Bank (COPDAB) Project. *Journal of Conflict Resolution* 24: 143-152.

Bremer, Stuart. 1996. Advancing the Scientific Study of War. In Stuart Bremer and Thomas Cusack (eds). *The Process of War: Advancing the Scientific Study of War*. Amsterdam, Netherlands: Gordon and Breach, pp. 1-33.

Bremer, Stuart and Thomas Cusack (eds.). 1996. *The Process of War: Advancing the Scientific Study of War*. Amsterdam, Netherlands: Gordon and Breach.

Craven, Mark, Dan DiPasquo, Dayne Freitag, Andrew McCallum, Tom Mitchell, Kamal Nigam, and Sean Slattery. 2000. Learning to Construct Knowledge Bases from the World Wide Web. *Artificial Intelligence* 118: 69-113.

Gerner, Deborah, Rajaa Abu-Jabr, Philip Schrodt, and Omur Yilmaz. 2002. Conflict and Mediation Event Observations (CAMEO): A New Event Data Framework for the Analysis of Foreign Policy Interactions. Paper prepared for the 2002 International Studies Association Convention, New Orleans, LA, March 2002.

Gerner, Deborah, Philip Schrodt, Ronald Francisco and Judith Weddle. 1994. The Analysis of Political Events Using Machine Coded Data. *International Studies Quarterly* 38: 91-119.

Jones, Daniel, Stuart Bremer, and J. David Singer. 1996. Militarized Interstate Disputes, 1816-1992: Rationale, Coding Rules, and Empirical Patterns. *Conflict Management and Peace Science* 15: 163-213.

KEDS. 2002. KEDS Data Sets: the Gulf. http://www.ukans.edu/~keds/data.html#gulf. Accessed January 29, 2002.

Mohr, Lawrence. 1982. *Explaining Organizational Behavior*. San Francisco: Jossey-Bass.

Pearl, Judea. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufman Publishers.

Putnam, Robert. 1988. Diplomacy and Domestic Politics: The Logic of Two-Level Games. *International Organization* 42: 427-460.

Schrodt, Philip and Deborah Gerner. 1994. Validity Assessment of a Machine-Coded Event Data Set for the Middle East, 1982-1992. *American Journal of Political Science* 38: 825-854.

Schrodt, Philip and Deborah Gerner. 1997. Empirical Indicators of Crisis Phase in the Middle East, 1982-1995. *Journal of Conflict Resolution* 41: 529-552.

Singer, J. David. 1965. Data Making in International Relations. *Behavioral Science* 10: 68-80.

Tomlinson, R.G. 1993. Word Event/Interaction Survey (WEIS) Coding Manual. Manuscript. Annapolis, MD. United States Naval Academy.

Vasquez, John (ed.). 2000. *What Do We Know About War?* Lanham, MD: Rowman and Littlefield.

Yang, Yiming, Jaime G. Carbonell, Ralf D. Brown, and Robert E. Frederking. 1997. Translingual Information Retrieval: Learning from Bilingual Corpora. *Proceedings of IJCAI 1997*.