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

Richard Stoll and Devika Subramanian
Rice University

February 6, 2002

1 Project Summary

With the end of the Cold War, the threat of a global-scale nuclear conflagration has receded. Unfortunately, this has not meant that all forms of serious conflict have been eliminated. Focusing only on the United States, we see that since the fall of the Berlin Wall, the United States has engaged in three wars (the Gulf War, Kosovo, and Afghanistan), as well as having its people - and even its territory - subject to attack. Other areas of the world have either experienced even more conflict (the Middle East) or threaten to do so in the immediate future (South Asia, the Korean Peninsula). Needless to say, should any of these conflicts break out and escalate, we might witness (albeit on a small scale) the nuclear conflagration that the US and the Soviet Union managed to avoid. We believe that the proliferation of news in electronic form as well as a series of advances in information extraction, data mining, statistical machine learning and stochastic modeling have made it possible to predict the outbreak of a serious international conflict by analyzing event data extracted from a multitude of sources over an extended period of time. The goal of our proposal is to develop techniques to construct extensive event data sets and models necessary to make such predictions.

We will build programs that will gather a large set of current and archived electronic information sources (e.g., Reuters, AP, UPI, cnn.com) and extract event data from them using established extraction techniques [GSFW94, $\text{CDF}^{+}99$]. Event data sets are coded time series where each element contains information on when an event occurs, who participated in it and what its characteristics are. We will develop algorithms for searching, organizing and visualizing large event data sets. Access to aggregated event-level data over a long period of time, can have a major impact on policymaking by providing an additional source of information upon which to base foreign policy decisions [Sch00]. We will develop new conflict prediction techniques that correlate event data streams across time and space (i.e., geographic regions). We will also learn stochastic process models from the available event data that track the evolution of conflict over time. We will assess the strengths and limitations of an objective, data-driven approach to modeling the genesis and evolution of conflict in various regions of the world. This project has the potential to change the science of conflict prediction in a fundamental way by exploiting the proliferation of information on political events in electronic form and by harnessing the speed and power of our computing engines to comprehensively and objectively analyze international conflicts.

2 Project Description

Our research will take advantage of the vast amount of information that exists (and is growing daily) on the World Wide Web. This information is in the form of websites maintained by media outlets throughout the world, such as Reuters. Many of these sites contain not only current issues, but extensive archives of older issues. We will use automated event extraction programs such as (`http://www.ukans.edu/keds`, [GSFW94])) to construct underlying event data from news reports. We will use time series of event data to generate predictions about the onset and progression of serious (militarized) international conflicts. Our preliminary analysis (discussed in detail below) leads us to believe that we will be able to predict the onset of these conflicts four to eight weeks in advance. The main thrust of our work will be

- to design information extraction techniques that build event data sets for use by the entire scientific community,
- to develop the algorithmic base for making predictions of serious conflict on the basis of event data, and
- to construct explanatory models of international conflict that builds on existing work in the area of international relations [Vas00].

In sum, our work will empirically test the information content of automatically extracted event data by predicting and building models of conflict directly based on them. It will extend existing work on automatic extraction of information from media sources and will integrate inductively learned models with known theories of international conflict.

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 [Sin65]. Events data – nominal or ordinal codes recording the interactions between international actors - are a prime example of this. Events data are widely used in quantitative international relations research. 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 (Conflict and Peace Data Bank) [Aza80] and WEIS (World Events Interaction Survey) [Tom93]) both cover only a limited period of time in the post-World War II era.

We will use existing techniques (Kansas Events Data Project [GSFW94],[CDF+99]) for automating the event extraction process and extend them as needed for the classes of sources we will consider for our work. 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 news magazines, 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 WEIS or COPDAB systems. The coding system specifies what types of political interactions constitute an "event," identifies the political actors that will be coded (for example, whether nonstate actors such as international organizations and guerrilla movements 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 machine-coding or automated information extraction, coding rules are implemented in systems such as KEDS or WebKB [CDF+99] 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 [CDF+99] relieving us of the burden of formulating such extraction rules manually. We will use the tools made available by WebKB [CDF+99] to rapidly construct special-purpose information extractors tuned to specific information sources.

KEDS [GSFW94, Sch00] has already been used to produce an extensive set of events data on the Middle East, 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 information extractors gatherers to more text sources on the World-Wide Web. In particular, we will 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 heirarchies as in [CYF+97] and use traditional corpora-based information extraction techniques on these sources. We also anticipate collecting events data from multiple English language sources. By comparing all these sources - foreign and English - we will build the most complete possible dataset of events. As well, we will be able to systematically evaluate the strengths and weaknesses of each source. It may be possible to eliminate the need to collect all these sources, but that is a decision that we will make after examining the data.

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. Organizing and indexing the data for effective retrieval, analysis and visualization will be important components of our work as shown in Figure 1.

2.2 Event Data Patterns and Early Warning

Given the collection of events data, we will move to predict, with a lead time of four to eight weeks, the outbreak of serious conflict. One particular form of serious international conflict that we seek to predict is militarized interstate disputes [JBS96]. 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[JBS96]). These represent the most widely used set of conflict cases that fall short of war. But as well, we 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

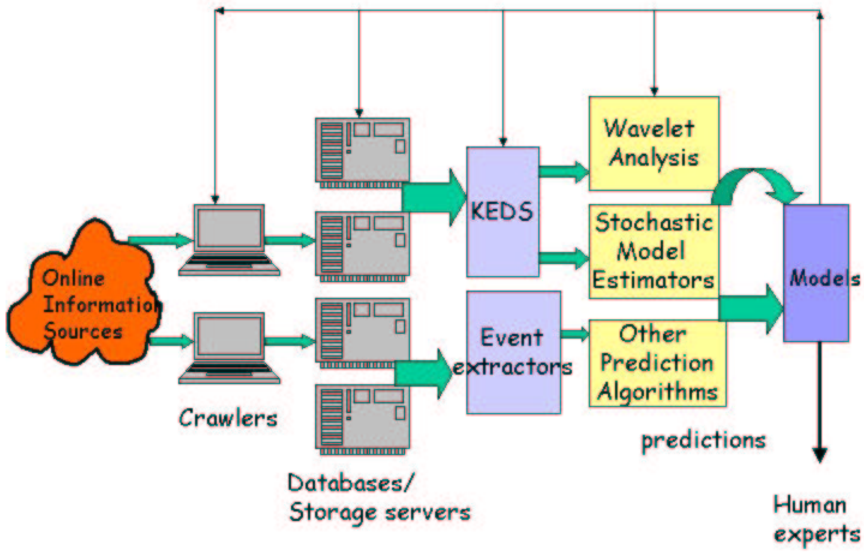

Figure 1: The architecture of the proposed system for extracting and analyzing event data from media sources around the world. Our project will push the boundaries of research in information extraction, model-based prediction and mining of event data. Note the feedback loops from predictions of models built from data back to the data gathering and analysis stages.

four to eight weeks prior to onset. An excerpt from an event data set from which we plan to make such predictions is shown in Figure 2.

The task of predicting the onset of a MID from a time series of scaled event scores is not easy. We illustrate the issues that arise in this context using a dataset involving events in the Middle East among twelve states from April 1979 to March 1999 (KEDS, 2002). Even a brief examination of this events data, shown in Figure 3 shows that it displays significant levels of non-stationarity. Conventional statistical and machine learning tools do not perform well on this data. We need to find the right granularity to aggregate the data in, and also to find the right window size to base prediction on. We aggregated event data in two-week periods, summarizing the total behavior in a single measure of the amount of conflict or cooperation experienced. This step averages out a lot of the noise in this time series. We then used 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, there is so much noise in the original time series, that its overall shape is not apparent upon visual inspection (see Figure 3). In a 256 level wavelet analysis, the trends in the data become very clear. At a scale value of 128, the resolution of the analysis decreases, but the trend plot (shown in Figure 3) clearly shows the ebb and flow of conflict in the region; and gives us confidence that there are extractable patterns in the data, and that prediction with a four to eight week horizon is possible using multiscale analysis of the time series of scaled event scores. The middle panel of Figure 3 shows successive approximations of the signal with progressively less "high-frequency" information. With the higher frequencies removed, what's left is the overall trend of the signal. At the higher frequency (small scales) we can perform a local analysis; at lower frequencies (large scale value) we can do global analysis. Combining local and global analysis is a useful feature of the wavelet method. The multiscale wavelet coefficients allow us

MID #	yr	mo	dy	c1	c2	init	targ	code	scaled score
3916	89	2	12	PAK	AFG	1	0	142	-1.1
3916	89	2	16	PAK	AFG	1	0	31	1
3916	89	2	16	AFG	PAK	0	1	32	1.9
3914	89	2	18	IRQ	IRN	1	0	31	1
3914	89	2	18	IRN	IRQ	0	1	31	1
3914	89	2	20	IRN	IRQ	0	1	66	1.9
3914	89	2	22	IRN	IRQ	0	1	201	-5
3914	89	2	23	IRQ	IRN	1	0	31	1
3916	89	2	24	AFG	PAK	0	1	111	-4
3914	89	2	27	IRQ	IRN	1	0	92	3.4
3916	89	3	4	PAK	AFG	1	0	94	-0.1
3916	89	3	12	PAK	AFG	1	0	223	-10
3916	89	3	20	AFG	PAK	0	1	23	-0.2

Figure 2: A portion of a constructed event data set (Middle East) for the first three months of 1989. Note the label identifying the event as part of a militarized interstate dispute (MID) taken from the MID database [JBS96].

to identify those periods in which there are significant levels of conflict. We can then build standard logit models that predict the presence or absence of high levels of conflict. Our predictor variables are summary measures of the level of previous conflict or cooperation.

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. Providing four to eight weeks advance notice of serious conflicts would be a significant achievement, but one that we feel is within our grasp.

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 understanding of the causes of conflict. To achieve this understanding, we have to 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, but most of it has involved predicting the onset of war. In recent years, there has been a convergence of findings, particularly those that account for the escalation of disputes to war; good summaries of the state of the discipline are in [Bre96, Vas00].

The factors that are most often associated with the escalation or deescalation of conflict - usually in situations involving a pair of countries - are ([Vas00],pg 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 that 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).

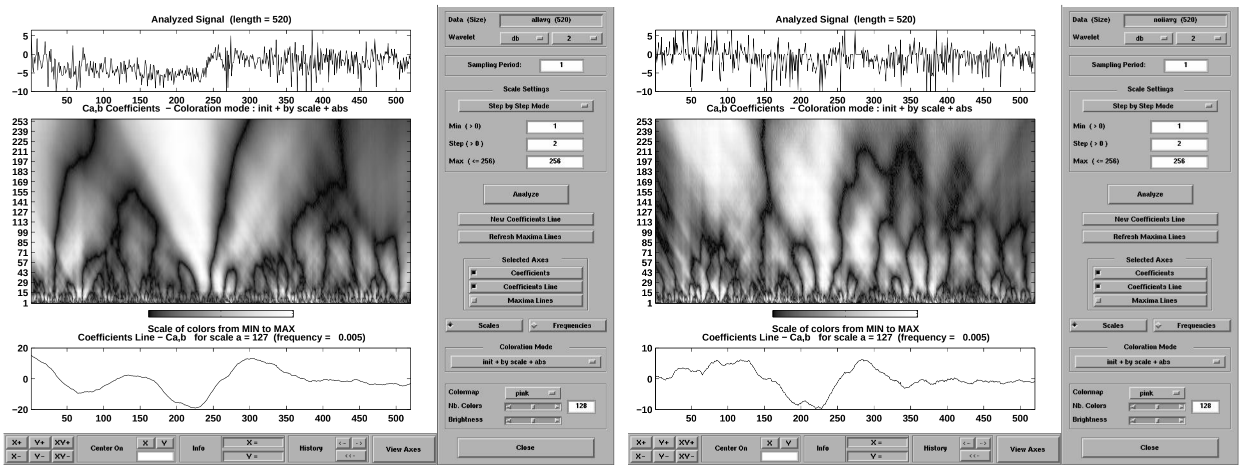

Figure 3: The figure on the left shows the analysis of the time series of scaled event scores averaged over all pairwise interactions among countries in the Middle East over two week blocks from April 1979 to March 1999. The figure on the right shows the same data with Iran and Iraq eliminated from the set. The Iran-Iraq war dominates events in this dataset and the figure on the right considers interactions with the Iran-Iraq war "subtracted" out of the data. The panel on the top shows the raw time series and the panel at the bottom shows a moving average over the signal (computed with a window of 128 weeks). The moving average plot highlights the non-stationarity in the data and is a summary of the ebb and flow of conflict in the region. The middle panel in both figures is a multi-resolution wavelet analysis with the Daubechies mother wavelet. The darkest regions in the panel correspond to significant singularities. In the left middle panel, the three lines that reach from the lowest scale to the highest correspond in time to the middle of the Iran-Iraq war, the Gulf War, and a brief war between Israel and Syria in 1994/1995. On the right, the middle panel highlights conflicts in the region involving countries other than Iran and Iraq – they become prominent once the effect of Iran and Iraq are subtracted from the analysis. The analysis was performed with the Continuous Wavelet 1-D tool in Matlab.

- Disputes involving states which 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; if across a series of disputes, if the parties take on more and more aggressive bargaining strategies, a war is 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 reduces 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. Otherwise, we have a set of interesting empirical relationships, but no real understanding of why conflict happens.

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 the 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 maybe unable to pursue the best choice for foreign behavior because this choice may not be supported by enough of the domestic groups. Second, leaders may take actions 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 [Moh82, Bre96]. The basis of explanation of a variance theory is causality. 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. A variance theory deals with variables, and efficient causes. Finally, 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 basis of explanation is probabilistic rearrangement. A precursor variable, X, is a necessary condition for the outcome. A process theory deals with discrete states and events. 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 menu 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 omelette.

We intend to treat the evolution of conflict as a stochastic process modeled by a dynamic belief network [Pea88]. A belief network is a representation of a probability distribution on a set of variables from the problem domain at hand. It comprises two parts: a qualitative representation and an associated quantitative one. The qualitative part of a belief network takes the form of an acyclic digraph. Each vertex in this graph represents a random variable that can take one of a set of values. The set of arcs of the digraph defines a dependence relation between the represented variables. Informally speaking, an arc is taken to represent a direct 'influential' relationship between the linked variables; absence of an arc between two vertices means that the corresponding variables do not influence each other directly. Associated with the qualitative part of the network is a numerical assessment of the strengths of the represented relationships as a set of (conditional) probabilities. These probabilities define a unique joint probability distribution on the variables discerned, that respects the independence relation portrayed by the graphical part of the network. Several algorithms for probabilistic inference are available [Pea88]. Dynamic belief networks (DBNs) are belief networks for temporal reasoning. A Markov process can be represented by two slices of a belief network representing the dynamics at time t and $t-1$. DBNs combine techniques from time series analysis and probabilistic reasoning to provide (1) a representation that integrates non-contemporaneous and contemporaneous dependencies and (2) methods for iteratively refining these dependencies in response to the effects of exogenous influences. Standard belief network inference can perform forecasting, control, and simulation on DBNs. The belief network formulation overcomes limitations in traditional time series analysis and allows us to move beyond the traditional assumptions of linearity in the relationships among time-dependent variables and of normality in their probability distributions. Recently there has been significant algorithmic interest in learning belief networks from data [FK02]. We intend to extend these methods to learn dynamic belief networks from data.

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

We have ambitious but feasible goals. First, we propose to build an extensive data archive of events between countries using the best tools and ideas in information extraction. Our objective is to build a robust data gathering system that updates event datasets 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 using wavelet analysis, and performing standard statistical and machine learning analyses that work on wavelet decompositions of the time series data, 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 nave 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 in the framework of dynamic belief networks. This will increase the knowledge base of the causes of conflict.

2.5 Educational Benefits

We believe that our research will have direct benefits for the classes we teach at both the undergraduate and graduate level. The algorithms for information extraction, fast wavelet analysis, learning predictive models based on wavelet coefficients, and bayesian network induction from data will be taught to the coming generation of majors in computer science. This will equip them with tools to model and solve prediction problems in a unique, real world application. A variety of courses in political science would also benefit from the research we propose. These courses range from a freshman level introduction to international relations, through a sophomore-junior level course on international conflict, to several undergraduate and graduate courses in international conflict. At the lower levels, the primary focus would be on charting the event stream, and integrating what it reveals into course lectures and discussions. At higher levels, students will learn the use of multi-resolution analysis methods like wavelets along with standard techniques in statistical analysis. Our primary focus is on the courses that we teach at Rice. We will make our course materials available for dissemination through our course web sites to faculty and students at other institutions.

2.6 Our Research and the Goals of the ITR Program

Our proposal on predicting and modeling conflict by extracting low-level event information over time from a variety of media sources

- provides better use of such information for societal use, including new ways of searching, organizing, preserving, and interacting with them.
- integrates data from multiple media sources and creates consistent data sets over vast scale in space and time.
- produces stochastic models of (aggregated) human interactions, enabling computational predictions of relatively rare, catastrophic, or revolutionary events (militarized interstate disputes).
- produces computational methods for learning stochastic dynamical system models (represented using dynamic belief networks) of international conflict from event data.

The research we outline in this proposal will accomplish these objectives. We will make data out of online media resources. These data can be used by a variety of different people to explore what has gone on in the world, as well as what is happening currently. These data will be analyzed to provide early warning indicators of the onset of serious conflict. This will involve a variety of methodological approaches, and we believe will yield important information about the onset of conflict that can be used by a wide variety of people. Finally, we will build and test models of the process by which conflict begins, escalates, and concludes. This will facilitate all of the goals noted above.

We believe that our proposal falls clearly in the guidelines for ITR submissions. It is research that is multidisciplinary. It takes advantage of the interfaces between computer science and political science. Finally, the data, methods, and results will be useful for researchers, policy makers, students, and others interested in matters of international events.

2.7 Results from Prior NSF Support

Dr. Subramanian's research on automatic reformulation was supported by grant IRI-89027121 from 1989 to 1993. It resulted in algorithms for designing new representations by explicit minimization of irrelevant computation. 3 Ph.D. students and over 15 undergraduates were trained under this grant. 3 edited collections of conference/symposium proceedings, a special issue of the Artificial Intelligence Journal, 6 journal publications and over 35 conference and workshop publications were produced under this grant. IRI-9319409 awarded to Dr. Subramanian from 1995 to the present funds her research on automating conceptual synthesis of opto-mechanical devices from specifications of behavior. It has resulted in real-time algorithms for conceptual design, and produced several novel designs for imaging systems of copiers as well as multi-legged walkers. 1 Ph.D. student and two undergraduates have been trained under this grant. 3 journal publications and 7 conference publications have been produced under this grant.

References

[Aza80] Edward E. Azar. The conflict and peace data bank (copdab) project. *Journal of Conflict Resolution*, 24(1):143–152, 1980.

[Bre96] Stuart Bremer. Advancing the scientific study of war. In Stuart Bremer and Thomas Cusack, editors, *The Process of War: Advancing the Scientific Study of War*, pages 1–33. Gordon and Breach,Amsterdam, Netherlands, 1996.

[CDF+99] M. Craven, D. DiPasquo, D. Freitag, A. McCallum, T. Mitchell, K. Nigam, and S. Slattery. Learning to construct knowledge bases from the world wide web. *Artificial Intelligence*, 1999.

[CYF+97] J. G. Carbonell, Y. Yang, R. E. Frederking, R. D. Brown, Y. Geng, and D. Lee. Translingual information retrieval. In *Proceedings of IJCAI 1997*, 1997. best paper award.

[FK02] N. Friedman and D. Koller. Being bayesian about network structure: A bayesian approach to structure discovery in bayesian networks. *Machine Learning*, 2002.

[GSFW94] Deborah Gerner, Philip Schrodt, Ronald Fransisco, and Judith Weddle. The analysis of political events using machine coded data. *International Studies Quarterly*, 38:91–119, 1994.

[JBS96] Daniel Jones, Stuart Bremer, and J. David Singer. Militarized interstate disputes, 1816-1992: Rationale, coding rules, and empirical patterns. *Conflict Management and Peace Science*, 15:163–213, 1996.

[Moh82] Lawrence Mohr. *Explaining Organizational Behavior*. Jossey-Bass, San Fransisco, 1982.

[Pea88] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Palo Alto, 1988.

[Sch00] P. A. Schrodt. Potentials and pitfalls in the application of event data to the study of international mediation. Technical Report Paper presented at the International Studies Association meetings, Los Angeles, 14-18 March 2000, Dept. of Political Science, University of Kansas, 2000.

[Sin65] J. David Singer. Data making in international relations. *Behavioral Science*, 10:68–80, 1965.

[Tom93] R. G. Tomlinson. World event/interaction survey (weis) coding manual, 1993. Manuscript, United States Naval Academy, Annapolis. MD.

[Vas00] John Vasquez. *What Do We Know About War?* Rowman and Littlefield, Lanham, MD, 2000.