Digging for Peace: Using Machine Learning Methods for Assessing International Conflict Databases

Robert Trappl, Johannes Fürnkranz, Johann Petrak

Austrian Research Institute for Artificial Intelligence Schottengasse 3, A-1010 Wien Email: {robert,juffi,johann}@ai.univie.ac.at

Technical Report oefai-tr-96-10

Abstract

In the last decade research in Machine Learning has developed a variety of powerful tools for inductive learning and data analysis. On the other hand, research in International Relations has developed a variety of different conflict databases that are mostly analyzed with classical statistical methods. As these databases are in general of a symbolic nature, they provide an interesting domain for application of Machine Learning algorithms. This paper gives a short overview of available conflict databases and subsequently concentrates on the application of machine learning methods for the analysis and interpretation of such databases.

1 Introduction

While enormous amounts of money have been and still are spent on the development of AI methods for military purposes, practically no effort is undertaken to use these methods to support the *prevention and termination* of conflicts and wars. The work presented in this paper is motivated by the deliberation that this area of research has not yet received the attention it deserves (Trappl 1986; Trappl 1992).

An important primary step is to understand the genesis and development of international crises as well as the success or failure of conflict management actions. Artificial Intelligence has lately been recognized as having some potential for supporting social scientists in this area of research (Mallery 1988; Hudson 1991; Unseld 1994; Schrodt 1996). Several approaches are possible:

- Understand the phenomenon 'conflict' itself and try to apply previously successful conflict resolution strategies to new or existing crises (Simpson Jr. 1985)
- Learn patterns in international events that lead to crises and use this information for early warning systems that can be used to timely alert peacekeeping organizations (Merritt, Muncaster, and Zinnes 1993)
- Detect regularities and rules that are shared between conflicts and use this information to gain insight in those parameters that support the escalation/deescalation of crises (Schrodt 1996; Mallery and Sherman 1993)
- Detect regularities and rules in conflict management actions and use this knowledge to increase the chance of success of subsequent mediation attempts (Bercovitch, Anagnoson, and Wille 1991; Bercovitch and Houston 1993; Bercovitch and Lamare 1993; Bercovitch and Langley 1993; Bercovitch and Wells 1993)

The work presented in this paper concentrates on the last two of these approaches: the analysis of databases of international conflicts with machine

learning algorithms. In section 2 we give an overview of some databases, and section 3 presents results of applying machine learning algorithms to the CONFMAN database of international mediation attempts.

2 Conflict Databases

One can distinguish between two primary types of conflict databases,

- *Event Databases:* describe the sequence of events that occur in crisis situations
- *Case-oriented Databases:* describe conflicts as a whole

Both kinds of databases have their advantages and drawbacks: Event databases do not need a rigid definition of what defines a conflict and how to specify the dates of outbreak and settlement, respectively. Also, they allow a more natural representation of what actually makes up a crisis situation. On the other hand it is often not clear which events are relevant to which conflicts. Case-oriented databases provide a clear documentation of information such as issues, fatalities, military power, and others of past conflicts but need a strict and often arbitrary definition of what should be considered to be a conflict.

There exist quite a few databases of either type, nearly all of them created for statistical analysis. Some examples are: the Correlates of War Militarized Interstate Disputes dataset (Gochman and Maoz 1984), the International Crisis Behavior (ICB) project (Brecher, Wilkenfeld, and Moser 1988; Wilkenfeld, Brecher, and Moser 1988), the COPDAB dataset (Azar 1980), the event data sets of the KEDS and PANDA projects (Schrodt and Davis 1994; Vogele 1994; Bond, Bennet, and Voegele 1994), the Butterworth dataset (Butterworth 1976), the KOSIMO database of conflicts (Pfetsch and Billing 1994), the CONFMAN database of mediation attempts (Bercovitch and Langley 1993), and the SHERFACS database (Sherman 1988).

We had access to three of the most comprehensive and well-known databases:

- KOSIMO (Pfetsch and Billing 1994): This database has been developed at the *Institute of Political Science* at the University of Heidelberg, Germany. One of the tables contained in this database describes 547 internal and international conflicts and wars between 1945 and 1990.
- SHERFACS (Sherman 1992): A large and complex database describing 1600 cases of quarrels and conflicts from 1943 to 1984 with roughly 5000 conflict phase descriptions.
- CONFMAN (Bercovitch and Langley 1993): A database of mediation attempts, developed under the supervision of Jacob Bercovitch from the Department of Political Science of the University of Canterbury, Christchurch, N.Z. The database contains descriptions of 921 mediation attempts from 241 disputes since 1945. Each dispute is described by 33 attributes, each mediation attempt by 12 additional attributes. The database contains several additional attributes that have been derived from the basic attributes or are still experimental. Nearly all attributes contain nominal values. The database has been previously used for statistical analysis of influence factors for successful mediation (Bercovitch, Anagnoson, and Wille 1991; Bercovitch and Houston 1993; Bercovitch and Lamare 1993; Bercovitch and Langley 1993; Bercovitch and Wells 1993).

3 A Case Study in Predicting Mediation Outcome

First we tried to learn decision trees for predicting the outcome of future conflict mediation attempts. The learning examples were taken from the CONFMAN database, but all entries where the outcome of the conflict management attempt is unknown were removed from the database. Furthermore we grouped the 5 different types of conflict management outcome into two classes: Mediation was *successful* when it resulted in a full or partial settlement of the conflict, or in a ceasefire. It was *unsuccessful* when a mediation attempt took place, but failed, or when mediation was only offered, but not accepted by the conflict parties. The resulting dataset consisted of 718 conflict management events, 408 (56.82%) of them resulting in failure and 310

(43.18%) being successful. Each event in this dataset was encoded with 52 attributes and one class variable that indicated whether the attempt has been successful or not.

As the basic induction algorithm we chose the standard decision-tree learning algorithm C4.5 (Quinlan 1993), because of its availability and flexibility. In a first attempt we generated and analyzed an unpruned tree with C4.5 (section 3.1). However, most of the nodes contained only a few examples, so that it seemed natural to generate simpler trees by pruning (section 3.2).

3.1 Analyzing an Unpruned Tree

The unpruned tree generated with -m 1 contains 547 nodes. Its accuracy on the training set is 99.7%, while its predictive accuracy (estimated with a 10-fold cross-validation) is about 60.3% compared to the default accuracy of 56.8%. The tree contains more than 100 leaves which split the training data into disjoint sets. Most of them contain only one example and cannot be expected to be predictive of the outcome of conflict management attempts. On the other hand, some leaves contain a relatively high number of examples that all had the same result. The unpruned tree contains 12 rules that contain 10 or more conflict management attempts. Five rules describe successful attempts, the other 7 cover failures.

Table 1 contains a summary of how many successful or unsuccessful conflict mediation events from how many different conflicts each of these rules describes. Together these 12 rules explain more than 25% (185 events) of the data set. Some of the rules are rather complicated, and it is unlikely that these regularities could have been detected by a human analyst. However, there are some simple rules testing only a few relevant conditions. For example rule **S1** says

> If there have been less than 400 fatalities and party B's raw power index is not extremely high and the conflict management type was mediation and the conflict lasted between 1 and 3 months
> then the conflict management was always successful in 15 mediation attempts in 8 different conflicts.

Rule	# Conditions	Success	Failure	# Conflicts
S1	5	15	0	8
S2	6	15	0	11
S3	10	14	0	5
S4	10	10	0	5
S5	12	12	0	7
F1	6	0	12	2
F2	9	0	19	4
F3	11	0	14	3
F4	12	0	16	5
F5	12	0	14	8
F6	8	0	13	3
F7	3	0	31	3
Total	_	66	119	_

Table 1: Rules that cover only successful (S1–S5) or only unsuccessful (F1–F7) conflict management attempts

On the other hand, rule **F1** shows us that

If there have been between 400 and 700,000 fatalities and party B's raw power index is not extremely high and both conflict parties have comparably high civil liberties and the conflict management type was mediation
then the conflict management was never successful in 12 mediation attempts in 2 different conflicts.

A complete listing of the interesting rules can be found in (Fürnkranz, Petrak, Trappl, and Bercovitch 1994). However, in general the tree contains too many leaves that cover only a very small number of examples. Therefore it seems natural to consider pruning heuristics to obtain simpler trees and rules.

3.2 Analyzing Pruned Trees

Table 2 gives an overview of some results we have achieved with different settings of two parameters of the standard decision-tree learning algorithm C4.5 (Quinlan 1993). For each setting we report the number of nodes (including leaves) in the generated tree (*Size*), the percentage of the training examples that will be correctly classified by the tree (*Purity*), and the predictive accuracy estimated by a 10-fold cross-validation and its standard deviation (Stone 1974).

Varying the -m parameter allows the user to constrain the tree generation by allowing only tests that have at least two outcomes with more than the specified number of examples. In particular this means that nodes that contain less than the specified number of examples will automatically become leaves and no further tests are considered. This prevents unreliable tests that are chosen near the leaves of the tree to discriminate small sets of examples from each other. Removing them leads to an increase in accuracy. However, a too high increase will cause the performance to decrease again, because C4.5 is forced to discard some relevant tests along with the irrelevant ones.

Varying the -c parameter on the other hand allows to specify the degree of pruning of the generated trees. Contrary to the minimum number of

Parameters	Tree Size	Purity	Predictive Accuracy
No Pruning (C4.5 -m1)	547	99.7%	$60.3\%~(\pm~4.8)$
C4.5 -m2	314	91.8%	$60.1\%~(\pm~3.3)$
C4.5 -m5	170	82.3%	$60.4\%~(\pm~5.7)$
C4.5 -m10	90	76.6%	$60.0\%~(\pm~5.2)$
C4.5 -m15	62	74.1%	$61.6\%~(\pm~4.7)$
C4.5 -m20	47	71.9%	$62.7\%~(\pm~2.0)$
C4.5 -m25	37	71.3%	$63.0\%~(\pm~2.2)$
C4.5 -m30	26	70.1%	$65.1\%~(\pm~2.5)$
C4.5 -m35	22	69.9%	$65.0\%~(\pm~4.2)$
C4.5 -m40	20	69.2%	$64.8\%~(\pm~2.6)$
C4.5 -m50	24	69.1%	$64.5\%~(\pm~3.5)$
C4.5 -c75	524	99.7%	$61.0\%~(\pm~4.5)$
C4.5 -c50	357	95.3%	$60.2\%~(\pm~3.6)$
C4.5 -c25	257	91.2%	$62.3\%~(\pm~4.4)$
C4.5 -c15	137	81.8%	$64.8\%~(\pm~4.6)$
C4.5 -c10	75	76.9%	$65.9\%~(\pm~4.9)$
C4.5 -c5	53	74.7%	$63.8\%~(\pm~6.0)$
C4.5 -c1	27	70.2%	$63.4\%~(\pm~5.8)$
C4.5 Default	173	86.2%	$62.5\% \ (\pm \ 5.2)$
C4.5 -m30 -c10	20	69.6%	$66.7\%~(\pm~3.7)$
Mode Prediction	1	56.8%	56.8%

Table 2: Decision tree learning results on the CONFMAN database.

examples criterion (-m, see above), pruning is a post-processing method that simplifies an existing tree¹ by replacing some of its internal nodes by leaves. The aim of pruning is the same as using -m (namely to discard unreliable nodes), but pruning is more flexible, because its parameter is independent from the actual number of examples used. Small values of the -c parameter cause more heavy pruning than large values.

We have also tried C4.5's default parameter setting (-m2 and -c25) as well as the combination of the best parameter settings (-m30 and -c10). The tree resulting from the latter settings (figure 1) has an estimated predictive accu-

¹The original trees in this series have been learned by using -m 1.

racy of 66.7% which is almost 10% above the accuracy of mode prediction.

The tree consists of 16 nodes², producing 9 simple rules. Each rule covers both, successful and unsuccessful conflict management attempts. Rule **P1** for instance is a generalization of rule **F7** from table 1 and illustrates that conflicts with a very high number of fatalities can hardly be solved by mediation. Obviously, conflict management activities have to be tried before too many fatalities occur. Rule **F7** contained two additional conditions that separated two successes and two failures from the 35 examples so that a cluster of 31 failures remained. However, it can be assumed that the two conditions that separate only four examples are irrelevant. Only three of the nine rules (**P3**, **P7** and **P9**) cover a majority of successful conflict management attempts.

3.3 Feature subset selection

Another method for obtaining simpler and more predictive trees by avoiding irrelevant distinctions of instances near the leaves of the tree is to limit the number of features that can be tested at the nodes of the tree. In Machine Learning this process is commonly known as *feature subset selection*. We have used the wrapper approach of (John, Kohavi, and Pfleger 1994) to determine the set of attributes from which the best decision tree can be learned. The algorithm starts with an empty set of attributes and greedily adds the attribute that gives the highest increase in estimated predictive accuracy for the tree that C4.5 grows from the new set of attributes. Alternatively, the algorithm can also choose to delete an existing attribute from the current set of attributes. Predictive accuracy is estimated with consecutive 10-fold cross-validation experiments (with different random splits) until the standard deviation of the resulting estimate is below 1%. If no feature can be added or deleted without decreasing the estimated accuracy of the tree for two consecutive tries, the program stops with the current set of features. In order to avoid to be too short-sighted a one-time decrease is not sufficient for

²In figure 1 four branches of the node *Management activity* have been turned into one single branch labelled *Other*. Hence it consists of only 16 nodes instead of 20 as specified in table 2.

stopping the algorithm. In this case two features may be added at a time if this increases accuracy.

Previous Relation of Mediator
Power Score and Disparity
Number of Involved Parties
Mediation Environment
Issues

Table 3: Relevant aspects for predicting mediation outcome

We have performed two experiments with different parameter settings for the basic induction module C4.5. Table 3 lists the relevant aspects of a conflict management attempt that are encoded in the variables judged important by both experiments³. In both cases the variable that describes the previous relationship of the mediator to the two conflict parties proved to be most important. Using a decision tree with only this one variable can raise the predictive accuracy from about 57% for always predicting the majority class to about 63%. Adding the next variable that in both cases reflects the power score of one of the conflict parties (although different parties are used in both cases) further increases the predictive accuracy to above 65%. These decision trees that test only two variables are already competitive with the best trees of table 2.

It is interesting to compare the results of feature subset selection with the results produced with classical statistical methods (Bercovitch and Lamare 1993) (see table 4). There is obviously a considerable overlap. Almost all of the variables of table 4 appear in one of the two experiments, most of them in both. The most notable exception is the absence of mediation strategy. The number of fatalities is also not among the most predictive features although if available it is very often chosen at the root of the trees. Obviously, C4.5's search heuristic based on information gain gives this attribute a high value, because one of its branches (fatalities $\geq 700,000$) is almost pure. However, the number of fatalities is partially reflected in the intensity of the conflict, which has been recognized as important, although only in one experiment.

³Details can be found in (Fürnkranz, Petrak, Trappl, and Bercovitch 1994).

Fatalities Mediation Environment Mediation Strategy Previous Relations of Mediator Issues Mediator Rank

Table 4: Relevant features for mediation outcome detected by statistical analysis

Using only fatalities for generating a decision tree would only yield 62.4% accuracy using the same grouping as in (Bercovitch and Lamare 1993). On the other hand, the Machine Learning method has attributed a higher significance to the previous relation of the mediator (63.3%). In addition, aspects concerning the power of the conflict parties and about the number of parties involved on each side have been considered.

4 Related Work

There have been several previous attempts to rule induction from international event databases (see (Mallery 1988; Schrodt 1991a; Schrodt 1996) for overviews).

(Schrodt 1991b) has performed similar experiments in predicting interstate conflict outcomes using the Butterworth "Interstate Security Conflicts, 1945–1974" (Butterworth 1976). He used his own implementation of ID3, the predecessor of C4.5, to learn decision trees for predicting the effects of management efforts with respect to five different outcomes. In all his experiments the estimated predictive accuracy of the learned trees was below mode prediction accuracy, i.e. below the accuracy that one would achieve by always predicting the majority class. However, his implementation of ID3 was not capable of dealing with numeric data and, more importantly, did not have C4.5's extensive pruning facilities. The only method used for getting simpler trees was manual feature subset selection, which did not result in gains in

accuracy. In our study, on the other hand, simple decision trees usually were able to achieve a higher predictive accuracy than an unpruned decision tree. However, even the unpruned tree obtained a significant gain in predictive accuracy compared to mode prediction. Predicting the outcome of conflict management attempts seems to be an easier task than to predict aspects of the outcome of the conflict itself. A reason for this might be that mediation events are more repetitive than the conflicts themselves.

(Unseld and Mallery 1993) have developed I²D, a variant of ID3 that was specifically developed to deal with the structured nature of the SHERFACS dataset (Sherman 1988). (Mallery and Sherman 1993) report a variety of rules that have been created by I²D. Again, the only simplification criterion was manual feature subset selection. This research focussed on learning single rules. The issue of predictive accuracy has not been addressed.

Situations of international conflict and war, like other complex human life situations, are often described and explained in terms of previous similar situations. Such comparisons often help to understand the various possibilities of actions the participants and international organizations can choose, and their possible consequences. Similarity-based case retrieval and analysis can therefore be a useful tool for analyzing a new conflict situation. Therefore, we applied case-based learning and similarity-based case retrieval methods to the KOSIMO database of conflicts (Petrak, Trappl, and Fürnkranz 1994). Figure 2 shows the retrieval of the five nearest neighbors of three selected cases in the database when using a similarity measure previously defined by a domain expert. The case "548 Bosnien-Herzegowina" has been coded and added to the database by one of the authors of the KOSIMO database for this experiment.

5 Conclusion

In this paper, we gave a short overview of databases of international conflict and conflict management actions and presented first steps of research on how inductive learning of rules with C4.5 can be used for the analysis of one of these databases, the CONFMAN dataset of mediation attempts. We plan to analyze a much larger and more recent version of the CONFMAN database in the future, using a wider range of machine learning techniques and including domain-specific background knowledge.

Acknowledgements

We are very grateful to Jacob Bercovitch, Peter Billing, John C. Mallery, Dwain Mefford, Frank Pfetsch, Philip A. Schrodt, Frank L. Sherman, and Sigrid Unseld for their many helpful comments and their continuous encouragement throughout our work.

This research was supported by the Austrian Federal Ministry for Science, Research, and the Arts.

References

- Azar, E. E. (1980). The conflict and peace data bank (COPDAB) project. Journal of Conflict Resolution 24(1), 143-152.
- Bercovitch, J., J. T. Anagnoson, and D. L. Wille (1991). Some conceptual issues and empirical trends in the study of successful mediation in international relations. *Journal of Peace Research* 28(1), 7–17.
- Bercovitch, J. and A. Houston (1993, October). Influence of mediation characteristics and behavior on the success of mediation in international relations. *Journal of Conflict Resolution* 4(4), 297–321.
- Bercovitch, J. and J. W. Lamare (1993). The process of international mediation: An analysis of the determinants of successful and unsuccessful outcomes. Australian Journal of Political Science 28, 290–305.
- Bercovitch, J. and J. Langley (1993, December). The nature of dispute and the effectiveness of international mediation. *Journal of Conflict Resolution* 37(4), 670–691.
- Bercovitch, J. and R. Wells (1993, January). Evaluating mediation strategies - a theoretical and empirical analysis. *Peace & Change 18*(1), 3-25.

- Bond, D., B. Bennet, and W. B. Voegele (1994). Panda: Interaction events data development using automated human coding. Extended Version of a paper presented at the 1994 Annual Meeting of the International Studies Association in Washington, DC on April 1st, 1994.
- Brecher, M., J. Wilkenfeld, and S. Moser (1988). Crises in the Twentieth Century - Handbook of International Crises, Volume I. Oxford: Pergamon Press.
- Butterworth, R. L. (1976). Managing Interstate Conflict, 1945-74: Data with Synopses. Pittsburgh: University of Pittsburgh Center for International Studies.
- Fürnkranz, J., J. Petrak, R. Trappl, and J. Bercovitch (1994). Machine learning methods for international conflict databases: A case study in predicting mediation outcome. Technical Report TR-94-33, Austrian Research Institute for Artificial Intelligence, Vienna.
- Gochman, C. S. and Z. Maoz (1984). Militarized interstate disputes 1816– 1976: Procedures, patterns, and insights. Journal of Conflict Resolution 28, 585–616.
- Hudson, V. M. (1991). Artificial Intelligence and International Politics. Boulder, CO: Westview Press.
- John, G. H., R. Kohavi, and K. Pfleger (1994). Irrelevant features and the subset selection problem. In *Machine Learning: Proceedings* of the Eleventh International Conference, Rutgers University, New Brunswick, NJ, pp. 121–129. Morgan Kaufmann Publishers, Inc.
- Mallery, J. C. (1988). Thinking about foreign policy: Finding an appropriate role for artificial intelligence computers. Master's thesis, M.I.T. Political Science Department, Cambridge, MA.
- Mallery, J. C. and F. L. Sherman (1993). Learning historical rules of major power intervention in the post-war international system. Paper prepared for presentation at the 1993 Annual Meeting of the International Studies Association.
- Merritt, R. L., R. G. Muncaster, and D. A. Zinnes (Eds.) (1993). International Event-Data Developments: DDIR Phase II. University of Michigan Press.

- Petrak, J., R. Trappl, and J. Fürnkranz (1994). The possible contribution of AI to the avoidance of crises and wars: Using CBR methods with the KOSIMO database of conflicts. Technical Report TR-94-32, Austrian Research Institute for Artificial Intelligence, Vienna.
- Pfetsch, F. R. and P. Billing (1994). Datenhandbuch nationaler und internationaler Konflikte. Baden-Baden: Nomos Verlagsgesellschaft.
- Quinlan, J. R. (1993). C4.5: Programs for Machine Learning. San Mateo, CA: Morgan Kaufmann.
- Schrodt, P. A. (1991a). Artificial intelligence and international relations: An overview. See Hudson (1991).
- Schrodt, P. A. (1991b). Classification of interstate conflict outcomes using a bootstrapped ID3 algorithm. *Political Analysis* (2).
- Schrodt, P. A. (1996). Patterns, Rules and Learning: Computational Models of International Behavior. University of Michigan Press. Forthcoming.
- Schrodt, P. A. and S. G. Davis (1994). Techniques and troubles in the machine coding of international event data. Dept. of Political Science, University of Kansas. Paper presented at the 1994 meeting of the International Studies Association, Washington DC.
- Sherman, F. L. (1988). Sherfacs: A new cross-paradigm, international conflict dataset. Paper written for presentation at the 1988 annual meeting of the International Studies Association.
- Sherman, F. L. (1992). Sherfacs research design, operational protocols, data representations and paradox 4.0 codebook. Project Report Number 3, Department of Political Science, Syracuse University, Syracuse, N.Y.
- Simpson Jr., R. L. (1985, June). A Computer Model of Case-Based Reasoning in Problem Solving: An Investigation in the Domain of Dispute Mediation. Git-ics-85/18, School of Information and Computer Science, Georgia Institute of Technology, Atlanta, Georgia.
- Stone, M. (1974). Cross-validatory choice and assessment of statistical predictions. Journal of the Royal Statistical Society B 36, 111-147.

- Trappl, R. (1986). Reducing international tension through artificial intelligence: A proposal for 3 projects. In R. Trappl (Ed.), Power, Autonomy, Utopia: New Approaches Toward Complex Systems. New York: Plenum.
- Trappl, R. (1992). The role of artificial intelligence in the avoidance of war. In R. Trappl (Ed.), *Cybernetics and Systems '92*, Singapore, pp. 1667–1672. World Scientific.
- Unseld, S. D. (1994, November). A selective overview of recent projects in Artificial Intelligence / International Relations. Paper prepared for the Second Workshop on the Potential Contribution of AI to the Avoidance of Crises and Wars, Vienna.
- Unseld, S. D. and J. C. Mallery (1993). Interaction detection in complex datamodels. MIT A.I. Memo No. 1298.
- Vogele, W. B. (1994). Global conflict profiles: Event data analysis using panda. Paper prepared for presentation at the annual meeting of the American Political Science Association, New York City, Sep. 1994.
- Wilkenfeld, J., M. Brecher, and S. Moser (1988). Crises in the Twentieth Century - Handbook of Foreign Policy Crises, Volume II. Oxford: Pergamon Press.



Figure 1: A decision tree generated with C4.5 -m30 -c10 from the CONF-MAN database.

****** Matching 548 Bosnien-Herzegowina 0.62 547 Deutschland-Tschechosl. (Muenchner Abkommen) 0.60 65 Israel I (Palaestinakrieg) 0.57 370 Zypern IV (Tuerkische Invasion) 0.55 289 Indien XVI (Kaschmir IV) 0.54 315 CSSR (Prager Fruehling) ****** Matching 547 Deutschland-Tschechosl. (Muenchner Abkommen) 0.77 315 CSSR (Prager Fruehling) 0.75 104 DDR (17. Juni 1953) 0.72 33 Griechenland (Buergerkrieg II) 0.67 52 Berlin I (Blockade) 0.66 219 Berlin III (Mauerbau) ****** Matching 477 USA-Grenada 0.66 281 Dominikanische Republik I (Intervention) 0.57 236 Kuba IV ('Kuba-Krise') 0.57 118 Guatemala I (Intervention) 0.57 352 Libyen-USA 0.57 14 Triest

Figure 2: The five nearest neighbors for three selected cases