The Possible Contribution of AI to the Avoidance of Crises and Wars: Using CBR Methods with the KOSIMO Data Base of Conflicts*

Technical Report OEFAI–94–32

Johann Petrak, Robert Trappl, and Johannes Fürnkranz
Austrian Research Institute for Artificial Intelligence
Schottengasse 3, A-1010 Wien
E-mail: {johann,robert,juffi}@ai.univie.ac.at

Abstract

This paper presents the application of Case-Based Reasoning methods to the KOSIMO data base of international conflicts. A Case-Based Reasoning tool - VIE-CBR - has been deveolped and used for the classification of various outcome variables, like political, military, and territorial outcome, solution modalities, and conflict intensity. In addition, the case retrieval algorithms are presented as an interactive, user-modifiable tool for intelligently searching the conflict data base for precedent cases.

*This work was supported by a grant of the Austrian Federal Ministry of Science and Research

Contents

1 Introduction

Situations of international conflict and war, like other complex human life situations, are often described and explained in terms of previous similar situations. For instance, the current situation in Bosnia is often analyzed by laymen as well as experts through comparing various aspects of the situation with Munich 1938 (appease "aggressor") on one hand, or Vietnam (send troops to aid "victim") on the other hand. Such comparisons often help to understand the various possibilities of actions the participants and international organizations can choose, and their possible consequences.

One assumption that lies behind this form of reasoning, is that similar situations often are best dealt with similarily. If there are differences, these might be valuable for adapting one's response to a situation or for pursueing an altogether different one.

Several formal methods that make use of this assumption have been developed in the context of statistics and AI. Case Based Reasoning (CBR) (Kolodner, 1993) is one of these methods. CBR has been used in various domains from military applications (Goodman, 1989), to reasoning about mediation strategies (Simpson Jr., 1985).

This paper presents a selection of CBR methods[1] we used (section 3) and their application to the KOSIMO[2] data base of international conflicts and wars. We chose the KOSIMO database, because it seemed to be the one of the available data bases that was best suited for CBR methods and because textual narratives (Pfetsch, 1991) are available for all cases in the data base. A short overview of the database is given in section 2; results from experiments with the database are presented in sections 4 and 5.

2 The KOSIMO Data Base

The KOSIMO database (Pfetsch & Billing, 1994) has been developed at the *Institute of Political Science* at the University of Heidelberg in Germany. It contains more than 1400 conflicts from 1482 to 1990. The database actually consists of three separate parts:

- A table "KRIEGE" of conflicts and wars dating from 1482 to 1945. This table contains almost no non-war conflicts. Also, for many cases information for several of the attributes is only partially available.

[1] the program package VIE-CBR (Petrak, 1994) is a Common-LISP program that implements all the methods descibed in this paper.

[2] KOSIMO is an acronym for **KO**nflikt**SI**mulations**MO**dell, i.e. conflict simulation model

- A table "PUTSCH" with over 400 putsches dating 1945 to 1990.
- A table "NOPUTSCH" with 547 internal and international conflicts and wars (but without putsches) dating 1945 to 1990, some still going on in 1990. This table contains not only all violent conflicts of the period (these are classified as "severe crisis" and "war"), but also a fair number of non-violent conflicts (classified as "latent conflict" and "crisis"), although these are somewhat underrepresented.

The NOPUTSCH table contains the most detailled and accurate information of these tables.

The database is an attempt to unify and extend case lists and databases of several previous research projects: primarily (Butterworth, 1976; Brecher *et al.*, 1988; Wilkenfeld *et al.*, 1988), but also (Gantzel & Meyer-Stamer, 1986), (Holsti, 1983), and others. There is also a thorough textual description of all cases in the NOPUTSCH table available in (Pfetsch, 1991). Moreover, (Pfetsch & Billing, 1994) contains various statistical analyses of the data. For these reasons, we chose to work with the NOPUTSCH table of the data base. In this paper, we will use the term KOSIMO database to refer to the NOPUTSCH table.

For a subset of the cases in the data base (i.e. the "basic conflicts" plus a few others, altogether 288 cases), additional fields were available in a seperate table. These fields contain structural information about the opponents, like difference in ideology, religion etc. This information was previously used in a statistical evaluation of these 288 cases (Billing, 1991).

In addition to the fact that the data base includes parts of previous work on conflict databases, the availability of the database and the cooperation of the two principal researchers, Frank Pfetsch and Peter Billing, were the primary reasons, why we chose to use that data base for our purposes.

2.1 Database Layout

The database was originally available as an SAS table, containing one record for each of the 547 cases. An example of a case from the SAS table is shown in figure 1. Each record contains about 70 fields. Most of these are used for coding information that does not directly describe the conflict (the identity of the coder, the sources used etc.), or were coded only for a limited number of cases in the database and therefore not used. 33 fields were finally extracted from the SAS table for further use. These fields are listed in table 1.

Fields in the SAS table are either numeric or contain text. Textual fields sometimes contain rather complex values. The field BETEILIG f.i. contains a list

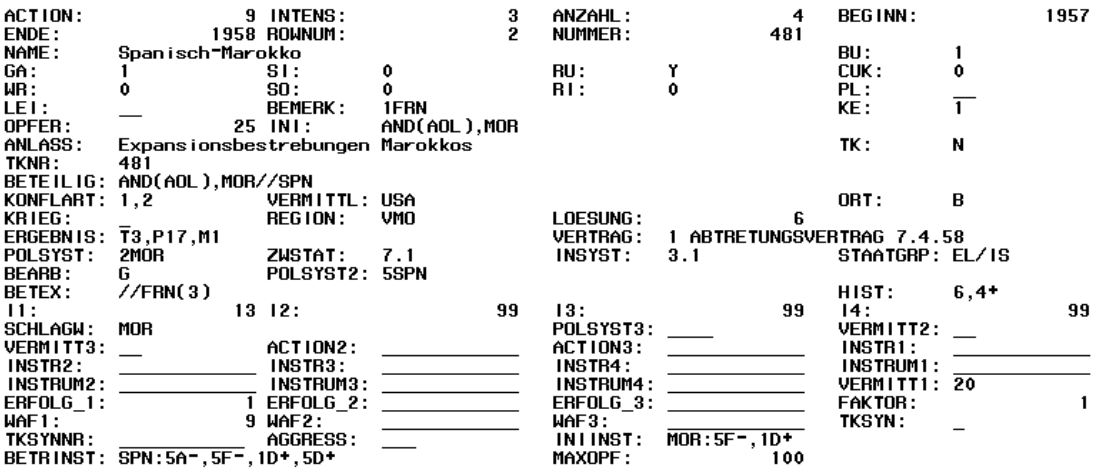

Figure 1: A KOSIMO case in the SAS table

Attribute	*Type*	*Description*
AGGRESS	text	aggressor (if different from initiator)
ANLASS	text	cause of conflict
ANZAHL	numeric	number of conflict parties
BEGINN	numeric	begin of conflict
BEMERK	text	remarks
BETEILIG	text	initiator and other parties
BETEX	text	external participants and their instruments
BETRINST	text	instruments used by the victim
ENDE	numeric	end of conflict
ERGEBNIS	text	political, territorial and military results
HIST	text	historical development
INI	text	initiator
INIINST	text	instruments used by the initiator
INSYST	text	reaction of the hegemonial powers
INTENS	numeric	conflict intensity
KONFLART	text	disputed goods and values
LOESUNG	numeric	modalities of solution
MAXOPF	numeric	fatalities (upper bound)
NAME	text	conflict name
NUMMER	numeric	conflict number
OPFER	numeric	fatalities (lower bound)
ORT	text	realm of influence of superpowers
POLSYST	text	political system of aggressor
POLSYST2	text	political system of victim
REGION	text	location of the conflict
STAATGRP	text	state of development of aggressor and victim
TK	text	partial conflict (diachron)
TKNR	text	number of partial conflicts (diachron)
TKSYN	text	partial conflict (synchron)
TKSYNNR	text	number of partial conflicts (synchron)
VERMITTL	text	mediator
VERTRAG	text	name of treaty
ZWSTAT	text	reactions of the neighbors

Table 1: Attributes extracted from the KOSIMO database

of three letter codes that identify the countries involved, separated by commas (see figure 1). Double slashes are used to indicate which states where involved on the side of the initiator of the conflict, and which on the target side. For some cases in the database even more than two sides are coded, using multiple double-slashes to separate them. In addition non-state participants are indicated by an expression of the format `AND(id)`, where `id` details the non-state organization.

Apart from a rather complex syntax, the encoding of field values in the KOSIMO databases has other properties that must be dealt with when converting to a more knowledge-based representation:

2.1.1 List-valued Fields

If a field contains a list of values, two interpretations are possible: the list might represent a set or multiset, or it might represent a sequence, in which case the order of the elements is significant. The KOSIMO database contains only set-valued fields.

Set-valued fields contain information, that – viewing the dataset as a relational data base – should be coded in a seperate table. An example would be the lists of instruments: to bring the database in first normal form, we would need to have one additional table for each of the fields BETRINST and INIINST (s. table 1). This table would have seperate rows for each combination of state/instrument that occurs for a specific case.

Conventional feature vector-oriented programs cannot easily deal with list-valued fields. Most ID3-like inductive learning programs, for instance, require attributes to have exactly one value (an exception is the I2D algorithm (Unseld & Mallery, 1993)).

With similarity-based algorithms, list-valued fields can be treated easily: As long as the concept of similarity is defined between any two possible elements, a similarity measure for sets of elements can easily be computed (s. section 3.1.1).

2.1.2 Multidimensional Fields

Some fields contain values that encode more than one conceptual entity (dimension). An example of such a field is LOESUNG (solution modalities): This field can contain one out of 11 possible values, that are used to encode (see table 2):

- whether the solution was aggreed on, forced upon one or both parties, or the result of a violent conflict
- whether a third party was involved to dictate the solution

- whether a conflict has ended or is still going on
- whether mediation was attempted

Value	*Description*
1	agreement
2	negotiated w/ 3rd party
3	authorative solution
4	negotiated solution forced by 3rd party
5	non-decision, ongoing conflict
5a	active retreat
6	policy of threat
7	ongoing violent conflict w/o mediation attempts
8	ongoing violent conflict w mediation attempts
9	ended violent conflict w/o mediation attempts
10	ended violent conflict w mediation attempts

Table 2: Possible values of the field LOESUNG (solution modalities)

This does not present a problem, if the possible values of such fields encompass all possible combinations of these dimensions. In this case, each dimension can be extracted without difficulty to a seperate field reserved for just that dimension. However, several fields in the KOSIMO data base contain only a limited number of combinations. To use the above example, the fact of attempted mediation is only stated for violent conflicts.

2.1.3 Semantic Dependencies

Semantic dependencies exist between fields, when a certain value of one field determines the value of another field. There are two possible kinds of dependencies:

- Functional dependencies: here one ore more fields determine through their values the contents of another field. With similarity based algorithms, functional dependencies can distort the actual weight of features.
- The dependent field does not make sense for certain values of the determining field. For instance, the field aggressor does not make sense, if no act of aggression occured. The correct way to deal with that problem would be to either use different case structures, or indicate the fact, that a field does not apply by a special reserved value. Unfortunately, the KOSIMO database instead uses the "empty value", which also indicates a "missing value".

2.1.4 Missing Values

The KOSIMO database does not employ a reserved value for indicating the fact, that information for a field could not be determined. Instead, the field is just left empty. This is no problem, as long as there is no other possible interpretation for the empty field. With set-valued fields, however, *two* interpretations are possible:

- The elements of the fields could not be determined
- The list is definitely empty

With the KOSIMO database, it is not possible to distinguish between these possible alternatives.

2.2 Conversion of the Database

We planned to use the information in the database for two purposes: experiments with inductive learning algorithms like C4.5 (Quinlan, 1993), and experiments with Case Based Reasoning methods. Since we planned to abstract from the exact case at hand, information about the parties involved was not of primary interest. For instance when converting the field INIINST (instruments of initiators) we did not keep the information *who* applied a certain instrument, but just created two lists of instruments used by either side of the conflict. Also, textual fields that cannot be automatically interpreted were removed (f.i. the field ANLASS, which contains a textual description of the direct cause of conflict). The field NAME (conflict name) was kept for easy case identification.

To convert textual fields with complex syntax (like INIINST) we defined the "language" of these entries by a context free grammar and converted the values to the intended format by creating a parser for that grammar[3]. The grammar for parsing the field INIINST is shown in figure 2. The automatic parsing of the fields INIINST, BETRINST, BETEILIG and BETEX also revealed a few errors in the original coding of these fields. These errors were subsequently corrected with the help of F. Pfetsch and P. Billing.

The list of fields finally included in the converted version of the KOSIMO database (case library NPNEW) is shown in table 3. Figure 3 shows the case from figure 1 after conversion for VIE-CBR.

In addition, a second case library was created from those 288 cases, for which structural information was available. The additional fields available in this library (case library NPSTR) are listed in table 4.

[3]This was accomplished by using a a YACC-like LALR parser generator, written in LISP by Mark Johnson

```
entries --> entry
entries --> entry ';' entries
entry   --> party
entry   --> party ':' insts
entry   --> party ':' ' ' insts
insts   --> inst
insts   --> inst sep insts
sep     --> ',' blanks
sep     --> blanks
bs      --> ' ' blanks
bs      -->
inst    --> digit digit anychar pm
inst    --> digit anychar pm
party   --> xyz
party   --> xyz ',' party
xyz     --> anychar anychar anychar
xyz     --> anychar anychar anychar '(' stuff ')'
stuff   --> anychar stuff
stuff   -->
pm      --> '+'
pm      --> '-'
```

Figure 2: Grammar for parsing the field INIINST

```
* (display-case c1 :long t)
OBSERVATION: 156        NAME: "Spanisch-Marokko"
BEGINN: 1957            ENDE: 1958         DAUER: 2
INTENS: 3               (ernste Krise)
SYSTEMEBENE: SE-INT     (internationalisiert)
POLSYST1: 2             (System im Uebergang)
POLSYST2: 5             (autoritaer)
GUETER: (KG-2 KG-1)     ((Kolonialbesitz Territorium))
EINFLUSSB: EB-B         (Ausserhalb Einflussb. der Grossmaechte)
OEKPOLNS: SG-NS         (Nord-Sued)
REGION: VMO             (vorderer/mittlerer Orient, Maghreb)
VERHHEGE: VH3D1         (A oder B maessigend auf K und L)
VERHZWSTAT: VZ7D1       (militaerische Konfr. K/L mit ausl. Unterstuetzung)
BETEILIGTE: (("AND" "MOR") ("SPN"))
EXTBETEILIG: (("FRN" 3))
ANZDIRBET: 3
INSTINI: (I5- I1+)      ((militaerisch/maessigend
                          bilaterale Diplomatie/verschaerfend))
INSTBETR: (I5- I1+ I5+) ((militaerisch/maessigend
                          bilaterale Diplomatie/verschaerfend
                          militaerisch/verschaerfend))
INSTEXTBET: (3)         ((dipl. Unterstuetzung))
ERGEBNISM: M1 (Waffenstillstand/Patt)
ERGEBNIST: T3 (Gebietserweiterung)
ERGEBNISP: (P17) ((Ziel teilw. erreicht (Kompromiss)))
LOESUNG: L-6 (Drohpolitik, Druck)
NUMMER: 481
DEMDIK: 1    INTENSP: 3   POLSTRUK: 0
MACHTDIFF: 1 IDEODIFF: 0  ENTWDIFF: 1
RELIDIFF: 1  KULTDIFF: 1  TRADDIFF: 1
SUMMDIFF: 5
```

Figure 3: A typical case in the library NPSTR

Field	Type	Usage	Description
name	text	descr	A short textual description of the conflict
observation	number	descr	The observation number of the case as stored in the original NOPUTSCH database
beginn, ende	number	descr	Year of beginning, and ending of the conflict
dauer	integer	out	The duration of the conflict, a calculated field derived from the fields beginn and ende
intens	even-range	out	One of four degrees of conflict intensity.
systemebene	symbol	in	The scope of the crisis: national, regional, international
polsyst1, -2	even-range	in	Type of political system for the conflict parties
einflussb	symbol	in	The scope of influence of this crisis
oekpolns:	symbol	in	The nord/south configuration of the conflict
region:	symbol	in	The geographical region where the conflict is located
anzdirbet:	even-range	in	The number of directly involved participants
instini	symbol-hierarchy,list	in	Instruments used by the conflict initiator
instbetr	symbol-hierarchy,set	in	Instruments used by the other conflict parties
instextbet	even-range	in	Instruments used by third party participants
verhhege	symbol-hierarchy	in	Actions of ,,hegemonial powers"
verhzwstat	symbol-hierarchy	in	Actions of involved parties
gueter	symbol-hierarchy list	in	Issues
ergebnism	symbol-range	out	Military outcome
ergebnist	symbol-range	out	Territorial outcome
ergebnisp	symbol-range list	out	Political outcome
loesung	symbol-hierarchy	out	Type of conflict resolution

Table 3: The fields converted to the case library

The final case libraries also contain two fields, that were not in the original database, but were calculated from existing fields: DAUER (duration of the conflict) and SYSTEMEBENE (scope of conflict). SYSTEMEBENE uses information from all the fields that contain direct or indirect participants and indicates, whether a conflict has international, regional or national scope.

Finally, two additional cases were added by a domain expert (P. Billing) to both libraries: "Bosnia 1992-" and "Munich 1938" (which are both not included in the original database because they are outside the covered timespan).

2.2.1 Knowledge About Field Values

The values of most fields in the database are either nominal or ordinal. To capture some of the relations between nominal values, we defined abstraction hierarchies for them. Abstraction hierarchies are useful when a similarity value between nominal values is needed: Without the additional information of the abstraction hierarchy, two values are either equal (best possible similarity), or not (no similarity). Using an abstraction hierarchy, one is able to assign to values that are related to each other through a common abstraction a similarity that lies within these extremes. An example for a abstraction hierarchy is shown in figure 4. Here, PS2 is similar to PS3, but not to PS5 (see section 3.1.1).

Field	Type	Usage	Description
machtdiff	symbol	in	Indication of difference in power
traddiff	symbol	in	Indication of difference in tradition
kultdiff	symbol	in	Indication of difference in culture
relidiff	symbol	in	Indication of difference in religion
ideodiff	symbol	in	Indication of difference in ideology
entwdiff	symbol	in	Indication of difference in economical development
polstruk	symbol	in	Indication of political structure
summdiff	even-range	in	Summury of differences

Table 4: The fields from the structure database

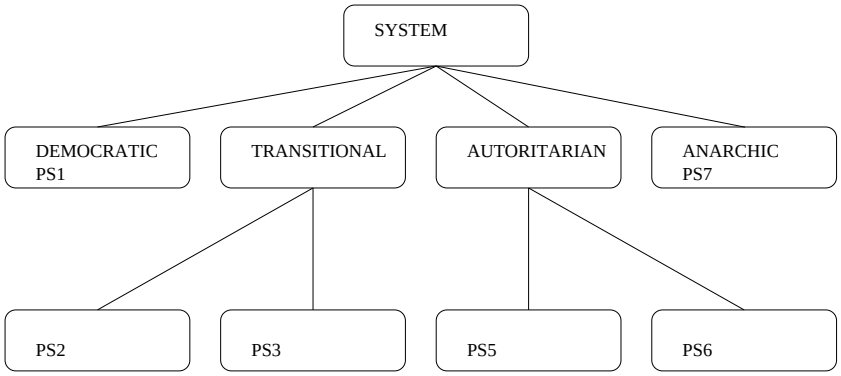

Figure 4: Abstraction hierarchy for field KONFLART (issues)

The abstraction hierarchies and ordering information for values used in the KOSIMO database was created with the help of the creators of the database. Tables 3 and 4 list the type of similarity information that was reconstructed for the fields in the case library (see section 3.1.1 for a detailled discussion of types of similarity measures).

3 Case Based Reasoning Methods

The term "Case Based Reasoning" is commonly used to refer to a group of algorithms, that are used for classification of problem solving purposes and have the following properties:

- a library of cases
- a method of finding "relevant" cases for a specific case at hand.
  - ◇ a method for adapting past solutions to a new situation
  - ◇ a method for evaluating a new solution
  - ◇ a method for storing new cases in the case library
  - ◇ a method for generalization of adaption rules based on that evaluation

Some of the methods indicated by ◇ might be present in only a rudimentary manner. Most case-based reasoning systems that are commercially available even lack them completely.

One form of CBR-algorithms are instance-based learning algorithms (Aha *et al.*, 1991) and nearest neighbor learning algorithms (Cost & Salzberg, 1993; Salzberg, 1991).

These methods do not use elaborate adaption techniques, but focus on complex similarity based case retrieval methods for performing a classification task.

3.1 VIE-CBR

VIE-CBR is an experimental system that shares most of the properties of commercially available CBR systems. However, it allows easy experimentation with alternative algorithms for any of its functional components. For the experiments below, version 1 of VIE-CBR was used (Petrak, 1994).

Similar to other programs, the primary representation format in VIE-CBR is a vector of features. However, features can contain any LISP data structure that has an external representation. Similarity measures can be defined as arbitrary LISP functions, but several functions for the most common feature types are predefined (see section 3.1.1 for a discussion of similarity functions used for the KOSIMO database). The primary method for case retrieval is a nearest-neighbor search, but arbitrary editing rules and k-nearest neighbor can also be used. The program also includes methods for estimating prediction accuracy by cross validation and leave-one-out testing. The following sections shortly describe the main components of the program.

3.1.1 Similarity Measures

Similarity is a concept that can be expressed in several different ways. It can be used as a binary predicate (,,A is/is not similar to B"), or it can be used as a fuzzy concept (,,A is not/somehow/very similar to B"). We will interpret similarity as a function that returns, for two objects A and B, a similarity measure in the range $0 \ldots 1$ where 1 means $A = B$ and 0 means there is no similarity between A and B.

Similarity[4] between cases is ultimately based on the similarities between field values. The first step in defining case similarities is therefore the definition of field similarities.

3.1.2 Field Similarity

The method used for determining the similarity between field values depends on the type of information and the semantics of the values in that field: for each type, a unique similarity function is created. The predefined similarity measures in VIE-CBR are explained below.

Even-Range: This similarity measure is intended for fields that contain a numeric value v out of a range of values $v_{min} \ldots v_{max}$. The values are assumed to be roughly evenly distributed within this range. The similarity of two values v_1 and v_2 is calculated as

$$\mathrm{sim}(v_1, v_2) = 1 - \frac{|v_1 - v_2|}{v_{max} - v_{min}} \qquad (1)$$

Symbol-Range: This similarity measure is intended for symbols out of a finited set N of symbols s_i, that can be arranged by some order relation into an ordered

[4] see also (Biberman, 1994; Richter, 1992)

sequence $< s_1, s_2, \ldots s_N >$. If $o(s_i)$ returns the rank of symbol s_i within the ordered sequence, the similarity of two symbols s_a and s_b can be calculated as

$$\mathrm{sim}(s_a, s_b) = 1 - \frac{|o(s_a) - o(s_b)|}{N - 1} \tag{2}$$

Symbol: This similarity measure is used, if there is no additional information for determining the similarity between values other than whether they are equal or not. Thus the similarity of two symbols s_a and s_b is

$$\mathrm{sim}(s_a, s_b) = \begin{cases} 1 & \text{if} s_a = s_b \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

(In (Cost & Salzberg, 1993) and (Biberman, 1994) alternative definitions for nominal scaled values are defined.)

Symbol-Hierarchy: Fields of this type also contain symbols, but the semantics of the symbols can be arranged in an abstraction hierarchy. Not every node in the abstraction hierarchy must actually be present in the data (in the KOSIMO data base, only the leaf nodes are present). Let s_{root} be the symbol at the root of the abstraction hierarchy, i.e. the most abstract symbol. Let for any two symbols s_1 and s_2 $msca(s_1, s_2)$ be the symbol that is the most specific common abstraction of these symbols. Let $a(s_1, s_2)$ be the number of arcs that must be traversed to get from node s_1 to node s_2 in the abstraction hierarchy. The similarity of symbols s_a and s_b is then calculated:

$$\mathrm{sim}(s_a, s_b) = 1 - \frac{(a(s_a, msca(s_a, s_b)) + a(s_b, msca(s_a, s_b))}{a(s_a, s_{root}) + a(s_b, s_{root})} \tag{4}$$

In VIE-CBR, the similarity definition for abstraction hierarchies requires the list of all arcs between nodes in the hierarchy, the hierarchy itself and all possible similarities is then calculated automatically.

Set of values: In addition, a field can contain a list of values of one of the above types (an example is shown in fig. 6). Lists are interpreted as sets, the order in which values appear is not significant. The similarity between two lists (sets) L_a and L_b is

$$\mathrm{sim}(L_a, L_b) = \frac{\mathrm{card}(L_a \cap L_b)}{\mathrm{card}(L_a \cup L_b)} + \sum_{\substack{e_i \in L_a - L_b \\ e_j \in L_b - L_a}} \mathrm{sim}(e_i, e_j) \tag{5}$$

Missing Values: In all similarity definitions, similarities involving a missing value are defined to be:

$$\begin{aligned} \mathrm{sim}(\texttt{nil}, \texttt{nil}) &= 0 \\ \mathrm{sim}(\texttt{nil}, x) &= 0 \end{aligned} \tag{6}$$

However, with the KOSIMO database, this definition cannot be used for set-valued fields, since no destinction is possible between an empty set and a missing value. An alternative definition with $\mathrm{sim}(\texttt{nil}, \texttt{nil}) = 1$ was therefore used. [5]

Figures 5 and 6 show some examples of similarity definitions for the KOSIMO data base.

```
(def-sim 'gut :symb-hierarchy
  '((kg)
    (kg-pol kg)(kg-mat kg)(kg-ord kg)(kg-int kg)(kg-8 kg)
    (kg-5 kg-pol)(kg-int kg-pol)
    (kg-1 kg-mat)(kg-7 kg-mat)
    (kg-3 kg-ord)(kg-4 kg-ord)
    (kg-2 kg-int)(kg-6 kg-int)))
```

Figure 5: Example of VIE-CBR similarity definition of type symbol-hierarchy.

```
(a)  (def-sim 'intens :even-range 1 4)
(b)  (def-sim 'einfluss :symb)
(c)  (def-sim 'gueter :list 'gut)
```

Figure 6: Examples of similarity definitions: (a) even range of symbols; (b) symbol; (c) list of symbols of type `gut`.

3.1.3 Case Similarity

In order to search the case library for cases that are similar to a case at hand, we need a method for determining an overall case similarity. There are several methods described in the literature.

For all experiments described in this paper, a weighted normalized sum was used. Let c_a and c_b denote two cases to be compared, F the set of attributes that are compared, v_f^a the value of attribute $f \in F$ in case a, v_f^b the value of attribute $f \in F$ in case b and w_f the wight of attribute f; the case similarity is then calculated as

$$\mathrm{csim}(c_a, c_b) = \sum_{f \in F} w_f \mathrm{sim}(v_r^a, v_r^b) / \sum_{f \in F} w_f \qquad (7)$$

[5] Another possibility would be to use an estimation of similarity that is based on the frequency of all possible values for the field.

Figure 7 shows an example of a case similarity definition.

In addition, the case similarity function might include a local weighting factor that indicates the "importance" of a case. If cases c_a and c_b have been assigned the weights w_a and w_b, respectively, equation 7 becomes:

$$\mathrm{csim}(c_a, c_b) = w_a w_b (\sum_{f \in F} w_f \mathrm{sim}(v_r^a, v_r^b) / \sum_{f \in F} w_f) \tag{8}$$

4 Similarity Based Classification

Several of the fields in the KOSIMO database describe aspects of the outcome of a conflict. Thus, they can be regarded as a classification of that case. The experiments described in this section tried to use similarity based case retrieval to estimate unknown outcomes.

The attributes of the case libraries were grouped into three categories:

- **descriptive**: attributes that are used only for case identification and informational purposes
- **input**: attributes that might be useful for inclusion into the a case similarity measure.
- **output**: attributes that describe the outcome of a case and should be predicted by the classification algorithm

Tables 3 and 4 show which fields were assigned to which category.

Several different experiments were carried out on either of the two case libraries (NPNEW and NPSTR), using various definitions of similarity measures, target variables, and prediction algorithms.

4.1 Estimation of Prediction Accuracy

In most experiments prediction accuracy was determined using a ten-fold *cross validation*. With this method, the original data set is divided into ten sets of cases of about the same size. Classification is then performed on every one of these ten sets, using every case of the used set as a case with unknown classifiaction. In each step the remaining nine sets are used as the case library that is searched for similar cases with known classification. After performing all ten steps of the cross

validation, the average prediction accuracy is used as an estimation of prediction accuracy on actual unseen cases [6].

We used two methods of calculating prediction accuracy:

- **Error Rate**: the relative frequency of incorrect guesses
- **Output Similarity**: the average similarity between guessed value and correct value. To calculate this measure, a similarity definition for the field containing the classification is required.

Calculating predicition accuracy using the output similarity method has the advantage giving different weights to "near misses" and totally wrong predictions. For instance, when predicting conflict intensity, it makes a difference whether a latent conflict is erroneously classified as crisis, or as war.

To estimate the actual "goodness" of prediction, class frequencies must be considered. For each classification experiment, we therefore also indicate the error rate / output similarity of an algorithm that always guesses the mode of the classification values (we will call this values "default error rate" and "default output similarity", respectively).

4.2 Similarity Definitions

The definition of case similarity measures was carried out with the help of domain experts (F. Pfetsch and P. Billing):

- SIM-EVEN: includes all input attributes with equal weight
- SIM-F: includes those attributes that were considered to be relevant by an expert, using weights estimated to indicated the relative relevance.
- SIM-PB: is based on statistic correspondences that were found in (Billing, 1991)
- SIM-PA and SIM-PC: are variations on SIM-PB with less attributes, that were estimated by an expert.

[6] Another method to estimate prediction accuracy is "leave one out" validation: here, one case after the other is picked for classification, using all the other cases in the case library for searching. This method was not used, because cross validation could directly be compared to inductive learning algorithms that we also used.

An additional similarity measure SIM-1 was used to analyze the contribution of the semantic information that was included into the similarity definitions. SIM-1 does not use any of these information, all values are treated as symbols that can only be equal or not equal.

All the case similarity definitions used in the experiments described in this paper, are shown in fig 7. All similarity definitions except SIM-OUT are input similarity definitions. All similarity definitions can be used with the NPSTR case library, all except SIM-PB can be used with the NPNEW case library.

4.3 Experiments

Table 5 lists the combinations of case library and similarity measure that were used in the experiments.

A summary of results is given in tables 6 and 7. Table 6 lists average output similarities and contrasts them with the default output similarities. table 7 lists errors rates and compares them with the default error rates.

For comparison C4.5, an inductive learning algorithm, was used for classification of all fields that are not list-valued. The learning algorithm was used to generate a decision tree and classification rules. C4.5 does not allow set-valued fields. One possibility to circumvent this limitation is to have a new field for each possible element in the set, that indicates, if the element is a member of the set. Since the number of possible different elements in each of the set-valued fields is rather large, we decided to use the abstraction hierarchies to limit the number of fields: one new field was generated for each node next to the root of the hierarchy; the field was set to true, if the set contains an element that is a leave of the subtree of that node, and to false otherwise.

Table 8 lists the error rates of classification with C4.5 on each of the libraries.

4.4 Results of Basic Retrieval Algorithms

Most of the prediction accuracies listed in tables 6 and 7 were only slightly better than a mode predictor would do. The fields that yielded the biggest difference in output similarity are LOESUNG, INTENS, and ERGEBNISM. The biggest improvements in error rate were achieved for ERGEBNISP and INTENS.

One should note that for these experiments, missing values were treated as just another value. Thus, all cases that had a missing classification but were classfied with a non-missing value, got output similarity zero. For the field ERGEBNIST, for instance, no value was specified in 63% of all cases in library NPSTR and 68% of all cases in library NPNEW.

Nr.	Library used	Similarity used
1	NPSTR	SIM-EVEN
2	NPSTR	SIM-F
3	NPSTR	SIM-PA
4	NPSTR	SIM-PB
5	NPNEW	SIM-EVEN
6	NPNEW	SIM-F
7	NPNEW	SIM-PC
8	NPNEW	SIM-1

Table 5: Overview of the experiments

Overall, the increase from default prediction accuracies is not dramatic. A possible explanation for this might be that the data is not well suited for classification tasks: The KOSIMO database looks at cases at a very high level of abstraction: the information contained in the variables of the database might not be the one required for prediction. One crucial information that is missing in the KOSIMO database is an indication of the sequence of events that make up a conflict.

Another problem is the absence of good adaptation rules (i.e. seperate domain knowledge that could be use to adapt classifications to certain properties of a situation). The fact that inductive learning with C4.5 did not do much better than VIE-CBR is an additional indication for this.

An interesting observation is, that in experiment 8 (where a similarity measure was used, that did not have any knowledge of abstraction hierarchies or symbol orderings) classification performance was very similar to those of experiments with other similarity measures. This might again be an indication for the inapropriateness of the database for classification purposes. It could also mean that hierarchy information was not relevant for classification or not be coded correctly (some of the hierarchies, f.i. the ones for the instrument codes, were difficult to code because of the multidimensionality of the values).

Figure 8 shows the distribution of input versus output similarities for classification of the field INTENS in experiment 1. When the library was searched for a similar case, there was always one that had at least a similarity of 0.62. There are four possible values of output similarity for the field INTENS. The four possible values of INTENS are not evenly distributed, therefore the distribution of output similarities favours the two highest values of similarity. Ideally, all output similarities should be 1, to indicate perfect classification.

Figures 9a and 9b show the frequencies of output similarities for field INTENS in experiment 1 for the actual classification algorithm (a) and a mode predictor

(b)[7]. A shift to higher output similarities clearly has occured.

4.5 Variations of the Algorithm

There are several enhancements of the basic retrieval algorithm described in the literature. We tested several of these variations with the KOSIMO database. It should be noted, however, that none of these algorithms is guaranteed (i.e. proven) to enhance the performance of the original algorithm.

4.5.1 k-Nearest Neighbor

Instead of retrieving just the single best match for a case, one could also retrieve the set of k best matches and use a voting scheme for selecting a classification from the cases in that set. The two voting schemes implemented in VIE-CBR are "majority" (Pick the class that occurs most often in the set) and "weighted majority" (Weight each occurrence with its similarity to the case to be classified. Pick the one with the highest weighted sum of occurences).

4.5.2 Case Weights

One might argue that, as far as classification is concerned, some cases are better suited as prototypes than others. A way to implement this thought in the classification algorithm is, to use classification performance on the *training* set as an indication of how good a case performs as a prototype. Before actual classification takes place on the test set, case weights are calculated for all cases in the training set. Case weights are calculated by dividing the number of times a case was used for correct classification by the number of times a case was used for classification. Cases that were never used for classification are assigned weight 0 (thereby removing them from the set of usable cases in the library). This is repeated a predefined number of times or until classification performance stops to increase on the training set.

4.5.3 Feature Weights

The most difficult task when defining case similarities is to decide which attributes might be relevant and how to weight their relative importance. There are several possibilities how feature weights might be determined automatically:

[7] Output similarity 0 never occured, since the mode of the field intensity (severe crisis) is not extremal. With a similarity definition of `even-range`, only a comparison between the opposite extremal values yields similarity 0.

Exp.Nr.	ERGEBNISM	ERGEBNIST	ERGEBNISP	LOESUNG	INTENS	Avg.
1	0.62/0.51	0.62/0.63	0.41/0.38	0.40/0.27	0.78/0.70	0.57/0.50
2	0.65/0.51	0.61/0.63	0.39/0.38	0.38/0.27	0.78/0.70	0.56/0.50
3	0.65/0.51	0.64/0.63	0.41/0.38	0.38/0.27	0.75/0.70	0.57/0.50
4	0.63/0.51	0.63/0.63	0.38/0.38	0.37/0.27	0.74/0.70	0.55/0.50
5	0.59/0.53	0.66/0.68	0.42/0.37	0.38/0.33	0.79/0.72	0.57/0.53
6	0.59/0.53	0.65/0.68	0.39/0.37	0.36/0.33	0.80/0.72	0.56/0.53
7	0.57/0.53	0.64/0.68	0.37/0.37	0.34/0.33	0.78/0.72	0.54/0.53
8	0.59/0.53	0.65/0.68	0.40/0.37	0.41/0.33	0.80/0.72	0.62/0.53

Table 6: Average output similarity (case based / predict mode).

1. Simulated annealing might be used to adjust randomly selected attribute weights and assess the adjustment by performing classification on the training set. Adjustments that increase the performance are kept. A similar technique is used in the nearest-neighbor algorithm of PEBLS (Cost & Salzberg, 1993)

2. For each classification step in the training set, analyze the similarities between each pair of corresponding attribute values. If the classification was correct, increase the weight of similar attributes and decrease the weights of differing attributes. If the classification was not correct perform the opposite adjustments.

3. Use some other method (f.i. statistical evaluation) for determining the relevant features.

VIE-CBR allows automatic adjustment of weights in a case similarity by method 2.

4.6 Results

Output similarities for the field INTENS when using some of the algorithm enhancements mentioned above are listed in table 9. None of the enhancements achieve a significant increase of classification performance. Performance for fields ERGEBNISM, ERGEBNIST, ERGEBNISP and LOESUNG also did not change much when compared to the original algorithm.

5 Similarity Based Case Retrieval

The classification process described in the previous chapters is based primarily on the process of "intelligent case retrieval". This process can also be valuable when

Exp.Nr.	ERGEBNISM	ERGEBNIST	ERGEBNISP	LOESUNG	INTENS
1	0.49/0.49	0.44/0.37	0.80/0.98	0.70/0.73	0.54/0.65
2	0.47/0.49	0.45/0.37	0.82/0.98	0.73/0.73	0.55/0.65
3	0.45/0.49	0.42/0.37	0.79/0.98	0.72/0.73	0.57/0.65
4	0.47/0.49	0.45/0.37	0.83/0.98	0.71/0.73	0.59/0.65
5	0.51/0.47	0.37/0.32	0.77/0.99	0.69/0.67	0.51/0.60
6	0.50/0.47	0.39/0.32	0.81/0.99	0.72/0.67	0.51/0.60
7	0.52/0.47	0.40/0.32	0.83/0.99	0.74/0.67	0.55/0.60
8	0.52/0.47	0.37/0.32	0.79/0.99	0.69/0.67	0.51 /0.60

Table 7: Average error rate (case based / predict mode).

Lib.	ERGEBNISM	ERGEBNIST	LOESUNG	INTENS
NPSTR (pruned tree)	0.43	0.36	0.74	0.53
NPSTR (rules)	0.46	0.39	0.66	0.51
NPNEW (pruned tree)	0.42	0.33	0.63	0.46
NPNEW (rules)	0.44	0.32	0.66	0.46

Table 8: Error rates achieved with C4.5

Exp.Nr.	None	Default	5-NN	5-NN Weighted	Case Weights
1	0.78	0.70	0.80	0.80	0.78
2	0.78	0.70	0.79	0.79	0.80
3	0.75	0.70	0.76	0.75	0.77
4	0.74	0.70	0.74	0.74	0.72
5	0.79	0.72	0.80	0.80	0.81
6	0.80	0.72	0.82	0.81	0.83
7	0.78	0.72	0.80	0.80	0.80
8	0.80	0.72	0.81	0.81	0.81

Table 9: Average output similarity for field INTENS for several algorithm enhancements.

systemebene	6	sysebene
polsyst1	3	polsys
polsyst2	3	polsys
einflussb	2	einfluss
oekpolns	2	nordsued
region	2	region
anzdirbet	6	anzbet
instini	3	inst-list
instbetr	2	inst-list
instextbet	1	extinst-list
verhhege	6	verhhege
verhzwstat	6	verhzw
gueter	6	gueter

(a) SIM-F

systemebene	1	sysebene
polsyst1	1	polsys
polsyst2	1	polsys
einflussb	1	einfluss
oekpolns	1	nordsued
region	1	region
anzdirbet	1	anzbet
instini	1	inst-list
instbetr	1	inst-list
instextbet	1	extinst-list
verhhege	1	verhhege
verhzwstat	1	verhzw
gueter	1	gueter

(b) SIM-EVEN

systemebene	1	symb
polsyst1	1	symb
polsyst2	1	symb
einflussb	1	symb
oekpolns	1	symb
region	1	symb
anzdirbet	1	anzbet
instini	1	symb-list
instbetr	1	symb-list
instextbet	1	symb-list
verhhege	1	symb
verhzwstat	1	symb
gueter	1	symb

(c) SIM-1

polsyst1	3	polsys
gueter	2	gueter
machtdiff	1	diff01
dauer	1	dauer
traddiff	1	diff01
kultdiff	1	diff01
relidiff	2	diff01
ideodiff	2	diff01
entwdiff	1	diff01
polstruk	1	diff01
summdiff	2	summdiff
region	2	region
verhzwstat	3	verhzw
verhhege	3	verhhege
einflussb	1	einfluss
anzdirbet	3	anzbet

(d) SIM-PB

gueter	2	gueter
dauer	1	dauer
region	2	region
verhzwstat	3	verhzw
verhhege	3	verhhege
einflussb	1	einfluss
anzdirbet	3	anzbet

(e) SIM-PC

polsyst1	3	polsys
gueter	4	gueter
summdiff	2	summdiff
region	2	region
verhzwstat	3	verhzw
verhhege	3	verhhege
anzdirbet	3	anzbet

(f) SIM-PA

Figure 7: The case similarities used for the experiments

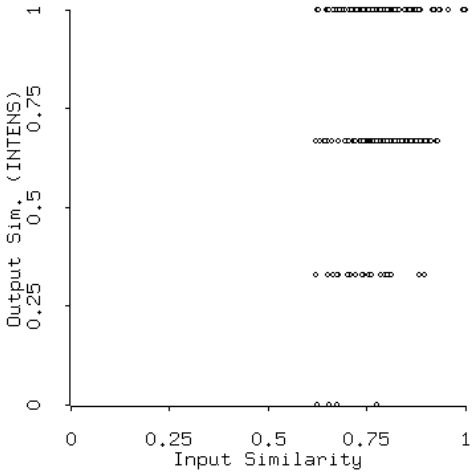

Figure 8: Input sim. vs. output sim. for field INTENS (Exp.1)

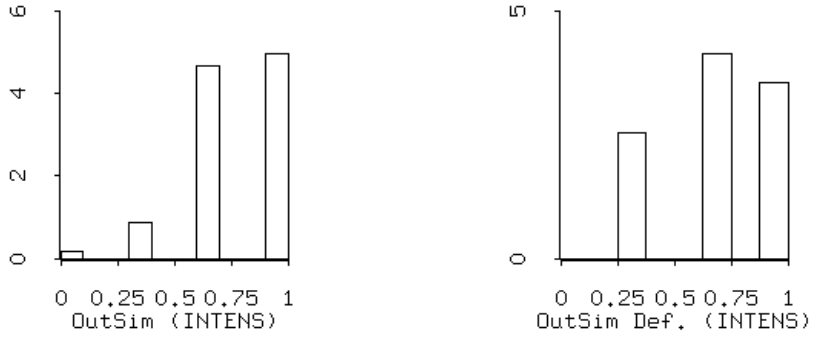

(a) VIE-CBR

(b) Mode Pred.

Figure 9: Frequencies of output similarities for field INTENS in experiment 1

used interactively by the user of the system. Intelligent case retrieval can aid an expert when he is trying to find precedent cases with a clearly defined concept of similarity. VIE-CBR allows experimatation with similarity measures and retrieval algorithms. In this chapter we present some results that were obtained by using VIE-CBR in this manner with the KOSIMO database.

When political scientists are asked to find precedent cases for a certain situation, they have no difficulty to find such cases, and justify the selection of these cases by listing several parallels between them and the case at hand. However, for different situations – and different experts – these parallels pertain to different features. VIE-CBR is useful, when one is interested in finding an "objective" list of similar cases.

5.1 Experiments

Each of the case similarity definitions SIM-PA, SIM-PB, SIM-EVEN and SIM-F were used for retrieval of similar cases. Fig. 10 shows the best 5 matches for the case "USA-Soviet Union 1960: U2 Incident" in the NPSTR library, using similarity definition SIM-PB.

A list of 10 cases that were selected by an domain expert (p. Billing) was used to retrieve the five best matches for each of them, using all of the similarity measures defined previously. The resulting list of matching cases was then given back to the expert. Most of the matched cases were found to be "somehow reasonable". Figure 11 shows the list of matching cases in library NPSTR using similarity measure SIM-PA.

5.2 Using Variations of the Algorithm

All of the algorithmic enhancements described in the previous chapter can also be used for interactive case retrieval. The procedure for calculating case weights can be used to analyze the suitability of cases as prototypes with respect to a certain set of classification variables and a certain similarity measure. Table 10 shows the stepwise increase in classification accuracy on the training set when case weights are calculated for classification of field INTENS in library NPSTR using input similarity SIM-EVEN. After 6 steps the values stabilize, having eliminated 115 cases from the original library of 289 cases. The distribution of case weights after step 6 is shown in figure 12.

```
* (list-sim-cases *s* 'sim-pb (list case130)
      :th 0 :max 5 :show-fields '(observation beginn name))
******* Matching 194 1960 USA-UdSSR (U2-Abschuss)
0.71 112 1954 UdSSR-USA (Chinesisches Meer, Piraterie)
0.72 320 1969 Schweden-UdSSR (Ostsee)
0.78 162 1958 UdSSR-USA (Sowjetischer Luftraum)
0.78 111 1954 UdSSR-USA (Luftzwischenfall)
0.91 195 1960 USA-UdSSR (RB-47 Abschuss)
```

Figure 10: The five best matches for case ,,1960: USA–UdSSR (U2)" using case similarity SIM-PB

Step	Inp.Sim.	Outp.Sim.	Error
1	0.80	0.78	0.54
2	0.74	0.82	0.44
3	0.69	0.82	0.41
4	0.66	0.84	0.40
5	0.66	0.86	0.36
6	0.65	0.86	0.37

Table 10: Classification on training set using case weight adjustment

```
******* Matching 548 Bosnien-Herzegowina
0.62 547 Deutschland-Tschechosl. (Muenchner Abkommen)  0.67
0.60  65 Israel I (Palaestinakrieg)                     1.00
0.57 370 ZYPERN IV (TUERKISCHE INVASION)                1.00
0.55 289 Indien XVI (Kaschmir IV)                       1.00
0.54 315 CSSR (Prager Fruehling)                        0.67
******* Matching 547 Deutschland-Tschechosl. (Muenchner Abkommen)
0.77 315 CSSR (Prager Fruehling)                        1.00
0.75 104 DDR (17. Juni 1953)                            0.67
0.72  33 Griechenland (Buergerkrieg II)                 0.67
0.67  52 Berlin I (Blockade)                            0.67
0.66 219 Berlin III (Mauerbau)                          0.67
******* Matching 477 USA-Grenada
0.66 281 DOMINIKANISCHE REPUBLIK I (INTERVENTION)       1.00
0.57 236 Kuba IV ('Kuba-Krise')                         1.00
0.57 118 Guatemala I (Intervention)                     1.00
0.57 352 LIBYEN-USA                                     1.00
0.57  14 Triest                                         0.67
******* Matching 363 Obervolta-Mali (Grenze I)
0.82 491 Burkina Faso - Mali (Grenze II)                1.00
0.80 253 Marokko-Algerien (Tindouf I)                   1.00
0.79 402 Sudan-Aethiopien                               1.00
0.77 389 Tunesien-Libyen                                0.33
0.77 459 Ecuador-Peru (Amazonas III)                    0.67
******* Matching 349 JEMEN AR-JEMEN VR
0.75  73 Syrien-Irak                                    0.67
0.69 418 Jemen VR-Jemen AR II                           1.00
0.69 370 ZYPERN IV (TUERKISCHE INVASION)                0.67
0.68   4 Griechenland (Buergerkrieg I)                  1.00
0.66  37 DOMINIKANISCHE REPUBLIK (INVASIONSVERSUCH I)   0.67
******* Matching 321 ARGENTINIEN-URUGUAY (RIO DE LA PLATA)
0.96 231 Bolivien-Chile (Lauca-Fluss)                   1.00
0.95 229 Brasilien-Paraguay (Parana)                    0.67
0.90 159 Argentinien-Chile (Palena-Disput)              0.67
0.90 459 Ecuador-Peru (Amazonas III)                    0.33
0.87  79 Saarland (Status)                              0.67
******* Matching 274 2. INDOCHINAKRIEG ('VIETNAMKRIEG')
0.86  90 Korea II (Korea-Krieg)                         1.00
0.70 257 Laos II (Buergerkrieg)                         1.00
0.69 130 2. INDOCHINAKRIEG (1.ABSCHNITT)                1.00
0.68 126 China (Tachen Inseln)                          0.67
0.68 172 China-Taiwan (Quemoy II)                       0.67
******* Matching 258 Malaya-Indonesien (Sarawak/Sabah)
0.74  58 Indien VII (Mahe)                              0.67
0.67 223 Indien XII (Goa II)                            0.67
0.67 378 USA-KAMBODSCHA (MAYAGUEZ)                      0.33
0.64 150 Aegypten-Frankr.,Grossbrit.,Israel (Suez)      1.00
0.63 345 Indien XVII (Bangladesh III)                   1.00
******* Matching 186 Italien (Suedtirol)
0.79 265 Indien XIII (Mizo)                             0.67
0.77 285 Kenia (Shifta-Ueberfall)                       0.33
0.77 183 CHINA (TIBET III)                              0.33
0.77 227 Irak (Kurden II)                               0.00
0.73 357 Afghanistan-Pakistan (Paschtunistan II)        0.67
******* Matching 131 Sudan (Autonomie im Sueden)
0.80 176 NICARAGUA (EXILANTEN)                          0.33
0.75 237 Jemen AR (Buergerkrieg II)                     1.00
0.75 306 Nigeria (Biafra-Sezession)                     1.00
0.75 386 ANGOLA (BUERGERKRIEG)                          1.00
0.73 403 ZAIRE (SHABA I)                                0.67
```

Figure 11: List of best matches for selected cases from library NPSTR, using similarity measy SIM-PA

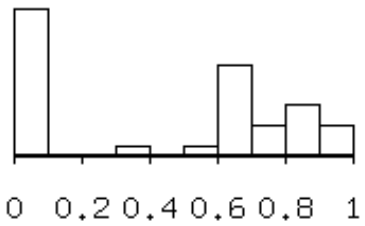

Figure 12: Distribution of case weights after step 6 (see table 10)

6 Conclusion

Working with the KOSIMO data, it became obvious that AI could contribute valuable methods to the study of conflict and war. There is, however a rather large initial effort necessary: the coding of international relations data using advanced knowledge representation techniques. It might be possible to automate parts of this process with the aid of natural language understanding systems (Alker Jr. *et al.*, 1991; Mallery, 1991).

Knowledge Representation

Knowledge representation techniques have been the subject of intensive research for as long as the field of AI exists. But little of that knowledge has actually been applied to the field of conflict research. Most data that has been collected in a machine-readable form was original intended to be used with statistical methods. Experience with the KOSIMO databases and other databases shows, that a rigid definition of coded concepts, the usage of structured representations instead of flat attribute-value representations could significantly enhance the value of a data set for more sophisticated methods. One step into this direction seems to be the SHERFACS database (Sherman, 1994). This data base uses a tree structure to describe the sequence of events taking place in each phase of a conflict.

Case Based Methods

Case-Based methods like those we have presented in this paper emphasize the importance of past cases for the classification and explanation of new cases. Thus, these methods correspond more closely than other formal methods to the way decision makers (that is, their advisory staff) tend to proceed in an actual crisis

situation. An important factor in such situations is the timely availability of relevant information. CBR methods might be able to play an important role in this context as "intelligent" data retrieval and interpretation tools. Another possible application of these methods might be in the environment of inter-governmental organizations where a timely reaction to early warning signals of conflict escalation is the most crucial factor for peacekeeping.

Any effort in this direction requires a tight cooperation of experts from each of the fields of applied AI, political science, and practical international relations. The experiences from the work presented in this paper has shown that especially the information retrieval component seems to be a promising field for future research.

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.

The support of the project "The Potential Contribution of Artificial Intelligence to the Avoidance of Crises and Wars" (principal investigator: Robert Trappl) by the Federal Ministry for Science and Research is greatly appreciated.

References

Aha, David W., Kibler, Dennis, & Albert, Marc K. 1991. Instance-Based Learning Algorithms. *Machine Learning*, **6**(1), 37–66.

Alker Jr., Hayward R., Duffy, Gavan, Hurwitz, Roger, & Mallery, John C. 1991. Text Modeling for International Politics: A Tourist's Guide to RELATUS. *In:* (Hudson, 1991).

Bergadano, Francesco, & Raedt, Luc De (eds). 1994. *Machine Learning: ECML-94*. Lecture Notes in Artificial Intelligence. Berlin: Springer-Verlag.

Biberman, Yoram. 1994. A Context Similarity Measure. *In:* (Bergadano & Raedt, 1994).

Billing, Peter. 1991. *Eskalation und Deeskalation internationaler Konflikte.* Politikwissenschaft XXXI, vol. 192. Frankfurt am Main: Peter Lang.

Brecher, Michael, Wilkenfeld, Jonathan, & Moser, Sheila. 1988. *Crises in the Twentieth Century - Handbook of International Crises.* Vol. I. Oxford: Pergamon Press.

Butterworth, Robert Lyle. 1976. *Managing Interstate Conflict, 1945-74: Data with Synopses.* Pittsburgh: University of Pittsburgh Center for International Studies.

Cost, Scott, & Salzberg, Steven. 1993. A Weighted Nearest Neighbor Algorithm for Learning with Symbolic Features. *Machine Learning*, **10**(1), 57–78.

Gantzel, K. J., & Meyer-Stamer, J. (eds). 1986. *Die Kriege nach dem zweiten Weltkrieg bis 1984.* München.

Goodman, M. 1989. CBR in Battle Planning. *Pages 264–269 of: Proceedings of a Workshop on Case-Based Reasoning.* Los Altos, CA: Morgan Kaufman.

Holsti, K. J. 1983. *International Politics: A Framework for Analysis.* 2nd edn. Englewood Cliffs.

Hudson, Valerie M. 1991. *Artificial Intelligence and International Politics.* Boulder, CO: Westview Press.

Kolodner, Janet. 1993. *Case-Based Reasoning.* San Mateo, CA: Morgan Kaufmann.

Mallery, John C. 1991. Semantic Content Analysis: A New Methodology for the RELATUS Natural Language Environment. *In:* (Hudson, 1991).

PETRAK, JOHANN. 1994. *VIE-CBR - Vienna Case-Based Reasoning Tool, Version 1.0: Programmer's and Installation Manual.* Technical Report TR-94-34. Austrian Research Institute for Artificial Intelligence, Vienna.

PFETSCH, FRANK R. (ed). 1991. *Konflikte seit 1945. Daten - Fakten - Hintergründe.* Freiburg: Verlag Ploetz. Five Volumes.

PFETSCH, FRANK R., & BILLING, PETER. 1994. *Datenhandbuch nationaler und internationaler Konflikte.* Baden-Baden: Nomos Verlagsgesellschaft.

QUINLAN, J. ROSS. 1993. *C4.5: Programs for Machine Learning.* San Mateo, California: Morgan Kaufmann.

RICHTER, MICHAEL M. 1992. *Classification and Learning of Similarity Measures.* SEKI Report SR-92-18. FB Informatik, Universität des Saarlandes, Saarbrücken.

SALZBERG, STEVEN. 1991. A Nearest Hyperrectangle Learning Method. *Machine Learning*, 251–276.

SHERMAN, FRANK L. 1994 (November). *SHERFACS - Research Design, Operational Protocols, Data Representations and Paradox 4.0 Database Codebook for SHERFACS 2.0.* Project Report Number 3.2. Department of Political Science, Syracuse University, Syracuse, N.Y.

SIMPSON JR., ROBERT 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.

UNSELD, SIGRID D., & MALLERY, JOHN C. 1993. *Interaction Detection in Complex Datamodels.* MIT A.I. Memo No. 1298.

WILKENFELD, JONATHAN, BRECHER, MICHAEL, & MOSER, SHEILA. 1988. *Crises in the Twentieth Century - Handbook of Foreign Policy Crises.* Vol. II. Oxford: Pergamon Press.