Efficient Search for Strong Partial Determinations

Stefan Kramer and Bernhard Pfahringer

Austrian Research Institute for Artificial Intelligence
Schottengasse 3
A-1010 Vienna, Austria
stefan@ai.univie.ac.at

Abstract

Our work offers both a solution to the problem of finding functional dependencies that are distorted by noise and to the open problem of efficiently finding strong (i.e., highly compressive) partial determinations per se. Briefly, we introduce a restricted form of search for partial determinations which is based on functional dependencies. Focusing attention on solely partial determinations derivable from overfitting functional dependencies enables efficient search for strong partial determinations. Furthermore, we generalize the compression-based measure for evaluating partial determinations to n-valued attributes. Applications to real-world data suggest that the restricted search indeed retrieves a subset of strong partial determinations in much shorter runtimes, thus showing the feasibility and usefulness of our approach.

Introduction

Functional dependencies (Mannila & Räihä, 1994) are a fundamental form of knowledge to be discovered in databases. In real-world databases, however, we have to face the effects of noise on functional dependencies: dependencies among attributes that would have been functional without noise are likely to have exceptions. Consequently, algorithms for inferring functional dependencies would not return those dependencies which are distorted by noise. More precisely and in machine learning terms, they would overfit the data, meaning that these algorithms would find too specific functional dependencies instead of the ones we would like to find.

In contrast to functional dependencies, *partial determinations* ((Russell, 1989), (Pfahringer & Kramer, 1995)) or *approximate functional dependencies* allow for exceptions. (In this respect, they are similar to *association rules* (Agrawal & Srikant, 1994).) Partial determinations may reflect probabilistic dependencies among attributes, but they may also be "impure" functional dependencies, i.e. functional dependencies which are distorted by noise and have only a few exceptions. In this paper, as in (Pfahringer & Kramer, 1995), we deal with those partial determinations which

help to compress a given database as much as possible. These highly compressive partial determinations will be called *strong partial determinations*.

In the next section we summarize the ideas from (Pfahringer & Kramer, 1995). Subsequently, we define and explain a new compression-based measure for partial determinations ranging over n-valued attributes. Then we describe an efficient method to search for strong partial determinations. In section 5, we report on experimental results of our method in several "real-world" databases.

Compression-Based Evaluation of Partial Determinations

Partial determinations are expressions of the form $X \rightarrow_d Y$, where X is the set of RHS (right-hand side) attributes, and Y is the set of LHS (left-hand side) attributes. d is the *determination factor* $d(X, Y)$ (Russell, 1989), the probability that two randomly chosen tuples have the same values of Y, provided they have the same values of X. In the following, we will restrict ourselves to RHSs consisting of single attributes.

Given a relation r, we define a mapping that corresponds to the expression $X \rightarrow_d Y$: the mapping relates the values of X occurring in r to the most frequently co-occurring values of Y. The tuples in r with Y values other than those majority values are called exceptions to the mapping in r.

The basic decision is which partial determinations to search for in databases. If we are interested only in accuracy, we are likely to get overly complex partial determinations in the presence of noise: we will find partial determinations fitting the noise instead of the underlying dependencies. To avoid this, we also have to take into account how complex partial determinations are. Therefore, (Pfahringer & Kramer, 1995) proposes a compression-based measure based on the so-called *Minimum Description Length* (MDL) principle (Rissanen, 1978). The MDL principle tries to measure both the simplicity and the accuracy of a particular theory (in our setting: a partial determination) in a common currency, namely in terms of the number of bits needed for encoding both a theory and the data
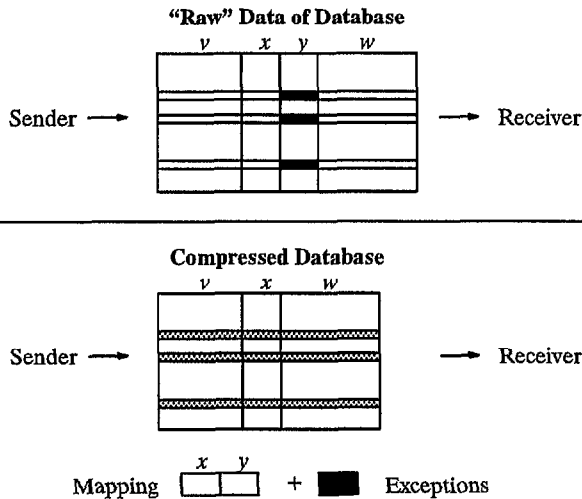
Figure 1: Transmission metaphor.

given that theory. The theory with the *minimal* total message length (the most compressive partial determination) is also the *most probable* theory explaining the data (Rissanen, 1978).

To illustrate the key idea of the measure for partial determinations proposed in (Pfahringer & Kramer, 1995), we consider the hypothetical task of transmitting a given database as efficiently as possible (see fig. 1). If we can find a good partial determination $X \rightarrow_d Y$ for a given attribute Y, transmitting the partial determination instead of the raw data may improve efficiency considerably: we just have to transmit the mapping (the *model*) and transmit the information to correct the values of Y for the exceptions to the mapping. Thus, we have to identify the exceptions (the "grey tuples" in figure 1) and transmit the corrections (the black box in figure 1). The values of Y need not be transmitted anymore. Thus we achieve some compression of the data, the degree of which is estimated by our measure based on the MDL principle.

The work reported in (Pfahringer & Kramer, 1995) has several limitations and problems: the paper only deals with partial determinations ranging over boolean attributes. Secondly, the problem of efficiently finding compressive partial determinations has not been solved. Thirdly, the method has not yet been applied to real-world databases. The following sections present extensions of our work that overcome these problems and limitations.

A New Compression-Based Measure for Multi-Valued Partial Determinations

In figure 2, we define the new compression-based measure for partial determinations ranging over multi-valued attributes. The total coding length (1) is the sum of the coding length of the mapping and the coding length of the corrections to be made. The coding length for a mapping (2) is the sum of the coding length

for specifying the LHS-attributes and of the coding length for the string of values for the RHS-attribute. The alphabet for this string is the domain of the RHS-attribute.

Exceptions to the mapping are encoded by (3). We have to identify the exceptions, and additionally encode a string of corrections for the identified exceptions to the mapping. Again, the alphabet of this string is the domain of the RHS-attribute.

The entropy-based bound of (Shannon & Weaver, 1949) is used as an estimate for the cost c_{choose} of encoding the selection of E elements out of N possible elements (4). Encoding a string over a multi-valued alphabet (5) is a straightforward generalization of (4). It can be thought of as a repeated binary selection encoding of what is the value occurring most frequently and which positions in the string are exceptions, i.e. have a different value. The resulting formula can be simplified to the definition (5) for c_{string}, which is just the sum of the entropy of the distribution of values in the string (6) and of an encoding of the values' identities sorted according to their frequencies (a kind of codebook for the particular string, 7).

Efficient Search for Strong Partial Determinations

Based on the view of partial determinations as functional dependencies which are distorted by noise, we propose an efficient search strategy to find strong partial determinations. Starting with overfitting functional dependencies, we propose to search for the best subset of left-hand side attributes according to our MDL-measure which avoids fitting the noise. In machine learning terminology again, we propose to *prune* the left-hand sides of overfitting functional dependencies. Pruning is achieved by a complete A*-like search over all subsets of LHS-attributes of the original functional dependency. In the following, we will refer to this two-level approach as *restricted search*. Limiting our attention to this highly interesting subset of partial determinations enables efficient search.

Even though both searching for functional dependencies and for partial determinations is exponential in the worst case, on average functional dependencies can be searched for much more efficiently due to much stronger pruning criteria. Additionally the evaluation function is much simpler: just counting the number of tuples in a projection vs. creating a mapping and computing its coding length for a projection. So in restricted search the expensive part has only to deal with a small number of attributes, namely those occurring in the left-hand side of the respective functional dependency. Furthermore, as a functional dependency is a partial determination with no exception, it also supplies an upper-bound on the coding length of possibly better partial determinations in the restricted search.

The next section will empirically show the effectiveness of restricted search in comparison to *full search*.

$$c(PD) = c(Mapping) + c(Exceptions) \tag{1}$$

$$c(Mapping) = c_{choose}(|UsedAttrs|, |AllAttrs|) + c_{string}(RHS\ Values\ Mapping) \tag{2}$$

$$c(Exceptions) = c_{choose}(|Exceptions|, |AllExamples|) + c_{string}(Corrections) \tag{3}$$

$$c_{choose}(E, N) = N * entropy(\{E, N - E\}) \tag{4}$$

$$c_{string}(String) = length(String) * entropy(char\_frequencies(String)) + c_{value\_order}(String) \tag{5}$$

$$entropy(\{F_1, \ldots, F_M\}) = -(\sum_{i=1}^{M} plog(F_i/N)),\text{where } N = \sum_{i=1}^{M} F_i \tag{6}$$

$$c_{value\_order}(String) = \sum_{i=0}^{j-1} log(M - i),\text{where } M \text{ is the cardinality of the alphabet,} \tag{7}$$

and j is the number of different chars in *String*.

$$plog(P) = \begin{cases} 0 & \text{if P=0} \\ P * log_2(P) & \text{otherwise} \end{cases}$$

Figure 2: The definition of the coding length c of a partial determination.

Full search is implemented as an iterative-deepening best-first search for partial determinations. For pragmatic reasons, if the number of attributes is larger than 20, both search approaches must be limited to left-hand sides of a prespecified maximum length, which is usually 10 in our experiments.

Empirical Results

For a preliminary empirical comparison of both search strategies we have done experiments using various small to medium-sized "real-world" databases. Breast cancer, Lymphography, and Mushroom are taken from the UCI-repository. Peace, Conflicts, and Attempts are databases capturing mediation attempts in international conflicts and crises (Trappl et al., 1996). Table 1 gives the sizes (number of attributes and number of tuples) of these databases and summarizes the search results. For the restricted search approach (pruning of overfitting functional dependencies) we list the number of attributes, for which a partial determination was found and the determinations' average compression factor. This compression factor for a given attribute is defined as the ratio between the coding length of the raw data and the coding length of the respective partial determination. Next we list the average compression factor achieved by a full search for partial determinations for the same RHS-attributes as those for which the restricted search found compressive partial determinations. Additionally we list the number of attributes for which *only the full search* was able to find compressive partial determinations, and again include their average compression factor. We can summarize as follows: for a subset of all attributes restricted search seems to find partial determinations which are on average almost as strong or compressive as those found by full search. Partial determinations missed by restricted search are on average significantly weaker (with the exception of the

Conflicts database). We have not yet performed detailed runtime measurements, but even for the smallest database Breast Cancer the difference is already of two orders of magnitude (80 seconds vs. 1 hour). For the other domains the differences in the runtimes were even more drastic. Summing up, the restricted search indeed retrieves a subset of strong partial determinations in much shorter runtimes.

Inspecting the retrieved determinations for useful knowledge also provided a few insights. For instance the Peace database includes both a few raw numerical attributes and their respective discretizations. Clearly a discretization should be functionally dependent on the raw data, which was *not* the case in the actual database. These erroneous discretizations were discovered as partial determinations. Moreover, we found strong partial determinations reflecting the implicit structure of the database: originally structured knowledge had been converted into one large and flat table due to restrictions imposed by standard propositional machine learning algorithms.

Related Work

(Shen, 1991) describes an algorithm that searches for determinations and two other kinds of regularities in a large knowledge base. Since the algorithm is just looking for determinations having a single left-hand side attribute, there is no need to avoid overfitting the data.

(Schlimmer, 1993) proposes an algorithm that returns every "reliable" partial determination with a complexity lower than a user-defined threshold. Since the evaluation function does not have a penalty for overly complex dependencies, it does not avoid overfitting

(Russell, 1989) and (Kivinen & Mannila, 1995) propose to determine the accuracy of partial determinations (respectively approximate functional dependen-

Domain	# A.	# T.	Restr.Search # PDs	Restr.Search Compression μ (σ)	Full Search Compression μ (σ)	Full Search # Add. PDs	Full Search Add. PDs Compression
Br. Cancer	11	699	6	1.49 (0.78)	1.64 (0.73)	5	1.20 (0.07)
Lymph.	19	148	14	1.13 (0.09)	1.13 (0.09)	5	1.01 (0.01)
Mushroom	23	8124	7	95.16 (26.84)	100.03 (23.25)	15	5.49 (6.32)
Peace	50	1753	33	4.47 (2.76)	4.63 (2.70)	16	2.11 (0.85)
Conflicts	33	268	24	1.38 (0.47)	1.38 (0.47)	4	1.37 (0.24)
Attempts	17	1753	10	2.23 (1.00)	2.26 (0.99)	7	1.65 (0.48)

Table 1: Results in real-world domains: the meaning of columns is explained in section 5.

cies) through sampling. However, if the proposed measures of accuracy would be used in a search framework, the returned determinations would overfit the data.

Conclusion and Further Work

Our contribution is twofold: firstly, we offer a solution to the problem of finding functional dependencies that are distorted by noise. Secondly, we propose a solution to the open problem of efficiently finding strong (i.e., highly compressive) partial determinations. Focusing attention on partial determinations derivable from overfitting functional dependencies enables efficient search for strong partial determinations.

Furthermore, we have generalize the compression-based measure for evaluating partial determinations to n-valued attributes. Applications to real-world data suggest that the restricted search indeed retrieves a subset of strong partial determinations in much shorter runtimes, thus showing the feasibility and usefulness of our approach.

Our approach to searching for partial determinations has several limitations: clearly, we cannot find all partial determinations, but only a subset. However, experiments indicate that those partial determinations which are distorted functional dependencies are among the most compressive dependencies to be found. It is also important to note that this approach is based on assumptions about real-world data, namely that there are strict functional dependencies which are distorted by noise. Thirdly, if the noise level is too high, no functional dependency might exist at all. So the pruning approach would not find partial determinations in these situations.

We plan to extend our work along the following lines: our method could be used to restructure the database in order to minimize the required memory, similar to INDEX (Flach, 1993). Another application of the measure could be to distinguish between reliable functional dependencies and those which are due to chance. Finally, it would be interesting to compare and combine our approach with sampling techniques as proposed in (Kivinen & Mannila, 1995).

References

Agrawal R. and Srikant R.: Fast Algorithms for Mining Association Rules. *Proceedings of the* 20^{th} *VLDB Conference*, Santiago, Chile, 1994.

Flach P.A.: Predicate Invention in Inductive Data Engineering, in Brazdil P.B.(ed.), *Machine Learning: ECML-93*, Springer, Berlin, pp.83-94, 1993.

Kivinen J., Mannila H.: Approximate dependency inference from relations, *Theoretical Computer Science*, 149(1):129-149, 1995.

Mannila H. and Räihä K.-J.: Algorithms for Inferring Functional Dependencies From Relations. *Data & Knowledge Engineering* 12 (1994) 83-99, 1994.

Pfahringer B., Kramer S.: Compression-Based Evaluation of Partial Determinations, *Proceedings of the First International Conference on Knowledge Discovery and Data Mining*, AAAI Press, 1995.

J. Rissanen: Modeling by Shortest Data Description. In: *Automatica*, 14:465-471, 1978.

Russell S.J.: *The Use of Knowledge in Analogy and Induction*. Pitman Publishing, London, 1989.

Schlimmer J.C.: Efficiently Inducing Determinations: A Complete and Systematic Search Algorithm that Uses Optimal Pruning. *Proceedings of the* 10^{th} *International Conference on Machine Learning*, 1993.

Shannon C.E. and Weaver W.: *The Mathematical Theory of Communication*, University of Illinois Press, 1949.

Shen W.-M.: Discovering Regularities from Large Knowledge Bases. *Proceedings of the* 8^{th} *International Workshop on Machine Learning*, 1991.

Trappl R., Fürnkranz J. and Petrak J.: Digging for Peace: Using Machine Learning Methods for Assessing International Conflict Databases. *Proceedings of the* 12^{th} *European Conference on Artificial Intelligence (ECAI-96)*, pp. 453-457, 1996.