

# Integrating Decision Support and Data Mining by Hierarchical Multi-Attribute Decision Models

Marko Bohanec<sup>1</sup>, Blaž Zupan<sup>2,3</sup>

<sup>1</sup> Jožef Stefan Institute, Jamova 39, SI-1000 Ljubljana, Slovenia  
marko.bohanec@ijs.si

<sup>2</sup> University of Ljubljana, Slovenia

<sup>3</sup> Baylor College of Medicine, Houston, U.S.A.

**Abstract.** An approach to the integration of Data Mining and Decision Support is proposed in which these two techniques are used jointly for solving evaluation or classification problems. The integration is based on a common modeling formalism: qualitative hierarchical multi-attribute decision models. These can be developed either (1) in a traditional Decision Support way through a dialogue between the decision maker and decision analyst, or by (2) Data Mining using a database of previously solved problems. It is shown that these methods can be combined in a number of useful ways: sequentially or in parallel, and with different levels of expert's involvement. The approach is illustrated using a real-world problem of housing loan allocation.

## 1 Introduction

Data Mining (DM) and Decision Support (DS) are two disciplines aimed at solving difficult practical problems. In many ways, they are complementary. To solve a particular problem, DS [13] tends to rely on knowledge acquired from experts, while DM [12] attempts to extract it from data. The integration of DM and DS may improve the quality of problem-solving methods, processes, and achieved results.

Several integrative approaches have already been suggested. Most often, they take a viewpoint of one discipline and try to enhance that discipline with methods of the other one. Thus, there are two general cases:

1. *Decision Support for Data Mining.* The idea is to enhance the DM process by supporting various decision-making tasks that take place in that process. For example, the approach based on ROC curves [16] supports the selection of the most appropriate machine learning algorithm for a particular DM problem.
2. *Data Mining for Decision Support.* Recent DS systems, which rely on technologies such as data warehouses, data cube and OLAP [14], incorporate more and more DM elements; a typical example is Microsoft's OLE DB for DM [15]. DM has also been used to improve multiple-criteria models in spatial decision-making [1].

In this paper, we propose a different and more "symmetric" approach: *using DM and DS jointly* for solving evaluation or classification problems. In DS, especially in Decision Analysis [10], such problems are traditionally approached by modeling: a

model is developed collaboratively by the decision maker and decision analyst, and used for the evaluation, classification and/or analysis of decision alternatives. Typical modeling techniques include decision trees, influence diagrams, and multi-attribute utility models. DM, on the other hand, operates on data about previously solved problems. Using a multitude of data analysis and machine learning methods, DM develops models that explain the data and solve previously unseen problems. In this case, models are typically represented by rules, decision trees, or neural networks.

The proposed integration of DS and DM is based on a common modeling formalism: qualitative hierarchical multi-attribute decision models. Such models are particularly convenient because there already exist both DS and DM methods for their development: DEX [4] for the development based on expert knowledge, and HINT [21] for the development from data. These two methods can be used independently or combined in a number of ways:

- *DEX only*: Applied when the expert is available, but no data for mining.
- *HINT only*: Used when there is a database of previously solved cases, but no expertise to guide model development.
- *Supervised*: This is a special HINT's mode that involves the expert, who can, at different levels and to different extent, guide the model development process.
- *Serial*: An initial model is developed by HINT from data and further refined by the expert using DEX.
- *Parallel*: Two or more models are developed in parallel by both HINT and DEX.
- *Combined*: One or more sub-models are developed by HINT, some others by DEX, and all combined together into a single model.

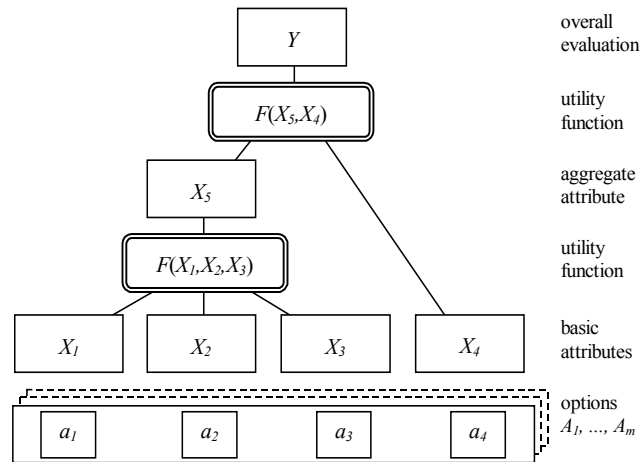
This paper is structured as follows. Section 2 provides a brief description of hierarchical multi-attribute models. Section 3 introduces a real-life decision problem of housing loan allocation, which is used as a working example. The methods DEX and HINT are presented in sections 4 and 5, respectively. Their integration is proposed in section 6. The paper is concluded by directions for further work and a summary.

## 2 Hierarchical Multi-Attribute Decision Models

*Hierarchical multi-attribute models* are widely used in decision analysis [3,10] for the classification or evaluation of options that are defined in attribute-value space. They represent a decomposition of a complex decision problem into smaller and less complex subproblems. In general, they consist of attributes  $X_i$  and utility functions  $F_i$  (Fig. 1). *Attributes* are variables that represent decision subproblems. They are organized hierarchically so that the attributes that occur on higher levels of the hierarchy depend on lower-level ones. We distinguish between *basic* attributes (terminal nodes) and *aggregate* attributes (internal nodes). Variables are connected by *utility functions* that aggregate partial subproblems into the overall evaluation or classification.

Common multi-attribute decision methods, such as AHP [18], deal with *quantitative* decision models. These are characterized by continuous attributes and utility functions that are typically defined in terms of attributes' weights. In contrast, the systems DEX and HINT, which are described later in sections 4 and 5, deal with

*qualitative* models. These are aimed at “soft” decision problems, i.e., less structured and less formalized problems that involve expert judgment. Qualitative models are characterized by discrete attributes and utility functions that are represented by decision rules (see section 3 for examples).



**Fig. 1.** Components of a multi-attribute model

The exploitation of models starts by representing decision *options* (alternatives) by the values of basic attributes: each option  $A$  is represented by a vector of values  $a_i$ . Options are then evaluated by a bottom-up aggregation, so that the overall evaluation ("utility") of each option is assigned to the root attribute ( $Y$  in Fig. 1). These values are used to compare and rank options, and eventually to select the best one.


### 3 A Real-World Case: Housing Loan Allocation

As a working example throughout this paper we use a decision model and data from a real-world management decision support system for allocating housing loans [5]. This system was developed for the Housing Fund of the Republic of Slovenia and represents one of the largest and most important applications of DEX. The Fund distributes financial resources to citizens in floats of loans. In each float, the task is to allocate available funds to applicants. Since 1991, the system has been used in 21 floats of loans, allocating the amount that corresponds to 254 million €. This is as much as two thirds of all housing loans approved in Slovenia in that period.

In each float, there are typically several thousands of applicants whose requested amount exceeds the available financial resources. Therefore, the applicants must be ranked into a priority order for the distribution of resources in accordance with the criteria prescribed in the tender. Each applicant is ranked into one of five priority classes. The criteria vary from tender to tender, but typically include:

1. applicant's *housing* conditions in terms of the ownership and suitability of present housing, the way of solving their problem, and the stage of solving;
2. applicant's *status* in terms of earnings, employment and the number of children;
3. *social* and *health* conditions.

In the system, the evaluation of loan priority is carried out by a qualitative hierarchical multi-attribute model whose structure reflects the above criteria (Fig. 2). In total, there are 12 basic and 7 aggregate attributes. The model was developed collaboratively by experts and decision analysts using DEX.



**Fig. 2.** Model structure for housing loan allocation

**Table 1.** Decision rules for the attribute *Soc-Health*

<i>Health</i>	<i>Social</i>	<i>Soc-Health</i>
(1) normal	(1) normal	(1) normal
(1) normal	(2) priority	(2) priority
(1) normal	(3) high_priority	(3) high_priority
(2) priority	(1) normal	(3) high_priority
(2) priority	(2) priority	(3) high_priority
(2) priority	(3) high_priority	(3) high_priority

The model is qualitative, so all the attributes are nominal: their values are words that typically express some level of priority. For example, there are three priority levels for *Social*: (1) normal, (2) priority, and (3) high\_priority. The specific value for each applicant is determined from his or her *Family Status* and *Age*, where the highest priority is typically granted for young families. The priority of *Social* is then combined with *Health*, which can be either (1) normal or (2) priority (for disabled people). The aggregation of *Soc-Health* is then defined by decision rules shown in Table 1.

#### 4 DS Method: DEX

The Housing Fund's model was developed using DEX [4], a computer program that facilitates the “DS way” of model development in collaboration between decision experts and analysts. The development is based exclusively on experts' knowledge and analysts' advice. DEX supports the following functions, which also roughly correspond to the main stages of decision model development and exploitation:

1. Acquisition of attributes and their hierarchy.
2. Acquisition and consistency checking of decision rules.
3. Description, evaluation and analysis of options.
4. Explanation of evaluation results.

During the last decade, DEX was used in more than fifty real-life decision problems [20]. Recently, there were applications in health-care [8], education [9], and in industry [7] for decision problems related to land-use planning, ecology, and the evaluation of enterprises, products, projects and investments.

## 5 DM Method: HINT

HINT is a data-mining method that develops hierarchical multi-attribute decision models using decision examples that may be taken either from an existing database of past decisions, or provided explicitly by the decision-maker. Each example is described by a set of attributes and its utility. The method is restricted to decision problems with nominal attributes and utility. Given an initial set of examples, the method develops a corresponding model in terms of a hierarchy of attributes and their definitions. The development proceeds either with or without human interaction. The method is described in detail in [21], and an extension of HINT that deals with noisy data is presented in [22].

**Table 2.** Set of examples that partially describes the function  $y=F(x_1,x_2,x_3)$ .

$x_1$	$x_2$	$x_3$	$y$
lo	lo	lo	lo
lo	lo	hi	lo
lo	med	lo	lo
lo	med	hi	med
lo	hi	lo	lo
lo	hi	hi	hi
med	med	lo	med
med	hi	lo	med
med	hi	hi	hi
hi	lo	lo	hi
hi	hi	lo	hi

HINT is based on function decomposition [2,11]. Let us describe the basic principles of operation using Table 2. It represents a set of examples that partially describe the utility function  $y=F(x_1,x_2,x_3)$ . From the viewpoint of hierarchical modeling, this situation corresponds to a “flat” multi-attribute model consisting of the root  $y$  and basic attributes  $x_1$ ,  $x_2$ , and  $x_3$ . The goal is to find a better hierarchical model in terms of structure and the complexity of utility functions. HINT does this by trying to decompose  $F$  into forms such as  $y=G(x_1,H(x_2,x_3))$ , where  $G$  and  $H$  are new utility functions.

The decomposition  $y=G(x_1,H(x_2,x_3))$  can be reformulated by introducing a new attribute  $c$  so that  $y=G(x_1,c)$  and  $c=H(x_2,x_3)$ . This represents a two-level hierarchy con-

taining  $c$  that depends on  $x_2$  and  $x_3$ . This indicates that HINT is in fact a concept learning method that induces a new variable (“concept”) in each decomposition step.

For the data from Table 2, all three possible non-trivial decompositions are shown in Fig. 3. The new tables are smaller than the initial one. The leftmost decomposition is the smallest and the easiest to interpret:  $c$  corresponds to  $\min(x_2, x_3)$  and  $y$  to  $\max(x_1, c)$ . Thus, this decomposition is preferred over the other two.

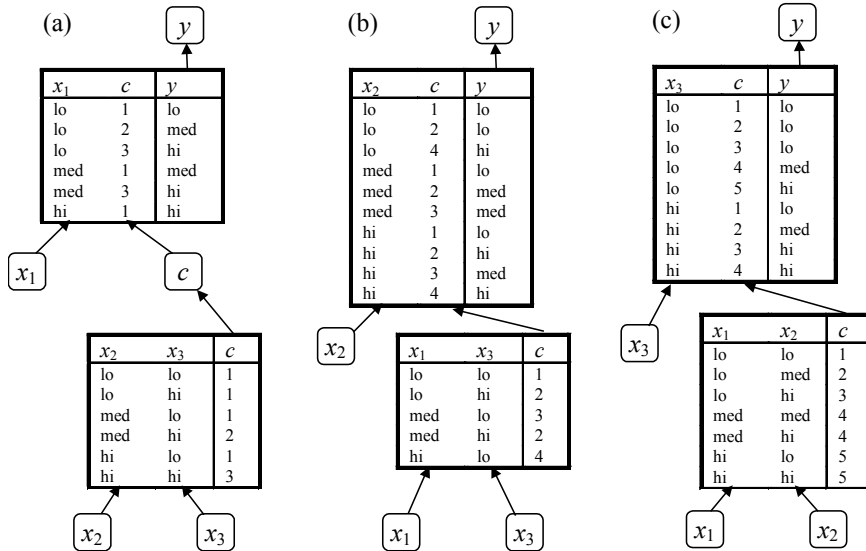


Fig. 3. Three different decompositions of Table 2

HINT repeats such decomposition steps recursively on  $G$  and  $H$ , as long as the complexity of the resulting model decreases. In general, the final result is a hierarchy consisting of the root  $y$ , basic attributes  $x_i$  (terminal nodes) and newly constructed concepts  $c_i$  (internal nodes). For each concept, HINT develops a corresponding utility function and represents it by sets of examples. The representation developed by HINT is therefore exactly the same as in DEX.

The decomposition algorithm itself is somewhat complex. In theory, optimal decomposition is a difficult combinatorial problem that consists of several NP-hard steps. For practice, HINT implements a number of shortcuts and heuristics. Instead of searching for all possible decompositions, it looks only for the functions  $H$  that have a small number of arguments, typically only two or three. The actual decomposition of  $F$  into  $G$  and  $H$  is carried out by a heuristic algorithm that is based on graph coloring, which is fast but sub-optimal. The complexity of decomposition is assessed by the number of values that are needed for the new concept  $c$ . For further details, see [21].

HINT was first implemented in C for several UNIX platforms, including HP-UX, SGI Iris, and SunOS. Recently, HINT became a component of a data mining framework called Orange [23], which uses Python for a scripting language and facilitates extensive experimentation and studies of different function decomposition scenarios.

## 6 Integrating DM and DS

In this section we show how DEX and HINT, the two representatives of DS and DM methods, can be combined in a number of interesting ways. We take the Housing Fund's decision model and data, and demonstrate six model development modes: DEX-only, HINT-only, supervised, serial, parallel, and combined.

### 6.1 DEX-Only Mode

In reality, the Housing Fund's loan allocation model was developed by a group of experts using DEX only. The development started in 1991, when the Fund was founded, and no data for mining was available at that time. The model was first applied in 1992, and then gradually adapted for subsequent floats of loans [6].

HINT became operational later, so it could not actually contribute to the project. Since then, however, the Fund has collected a considerable amount of data about loan allocation, which we use to experimentally assess the hypothetical contribution of HINT had it been available earlier. This study is presented in the following sections.

### 6.2 HINT-Only Mode

For the experimental evaluation of HINT, we took applicants' data from one of the floats carried out in 1994. There were 1932 applicants in that float. Each data item contained 12 two to five-valued basic attributes. Due to the discreteness of attributes, the 1932 records provided 722 unique examples, covering 3.7 % of attribute space.

The first goal was to *reconstruct* the original decision model using only the available applicants' data, supplemented by the already known loan priority. For this purpose, each example was classified by the original evaluation model and the resulting unstructured database was submitted to the decomposition method. HINT was used in the so-called *unsupervised* mode, which proceeds without human intervention.

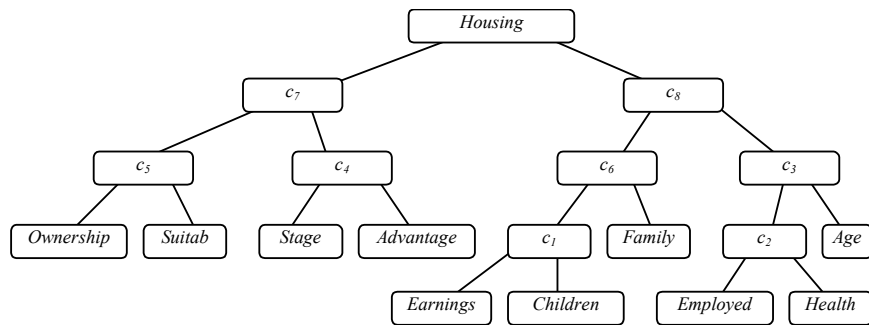


Fig. 4. Model structure developed in HINT-only mode

HINT developed the model whose structure is shown in Fig. 4. The comparison with the original model (Fig. 2) reveals there are two missing attributes: *Cult-Hist* and *Fin-Sources*. This is because these attributes affect the loan priority only under some rare special circumstances. For example, *Cult-Hist* applies only to houses that are declared a cultural or historical monument. There were no such cases in the database.

Further comparison reveals that in some parts the models match very well. This is especially true for the concepts  $c_1$  and  $c_5$ , which closely correspond to the original attributes *Status* and *Present*, respectively. When reviewed by a domain expert, some other parts were found less satisfactory and more difficult to interpret, especially around  $c_2$  and  $c_3$ . At this point, the expert clearly expressed the need for supervised decomposition (section 6.3).

With respect to generalization, HINT performed very well. Its classification accuracy, assessed by a 10-fold cross validation, was  $94.7 \pm 2.5$  %. By this, HINT outperformed C4.5, a state-of-the-art machine learning program that induces decision trees from examples [17], which achieved the accuracy of  $88.9 \pm 3.9$  %.

### 6.3 Supervised Mode

Central to each decomposition step is the partition of attributes  $X$  to  $A$  and  $B$ , so that  $F(X)=G(A,H(B))$ . In the *unsupervised* mode, partition selection is guided by HINT assessing the complexity of  $G$  and  $H$ . In the *supervised* mode, the user is allowed to intervene by suggesting the most suitable partition; a few best partitions are presented by HINT to the user, who selects one and labels the newly constructed concept.

The supervised mode effectively combines the DM and DS principles of model development. The process basically proceeds in the DM way from data, but can be complemented with expert's knowledge in the DS sense.


To illustrate this process, let us use the same data as in section 6.2. After removing the irrelevant basic attributes *Cult-Hist* and *Fin-Sources*, the data set contains 722 learning examples described by 10 basic attributes; each example has a known loan allocation priority (*Housing*). This situation corresponds to a flat decision model.

The data set is then examined for decomposition: all possible partitions of the ten input attributes with  $B$  containing 2 or 3 attributes are examined. From these, HINT selects only those in which the newly created attribute  $c$  has the lowest number of discrete values. In our case, the lowest number of  $c$ 's values was 3, and among 120 possible partitions there were 11 selected and presented to the expert:

- |                                      |                                    |
|--------------------------------------|------------------------------------|
| 1. <i>Suitab Advantage Employed</i>  | 7. <i>Earnings Employed Family</i> |
| 2. <i>Advantage Stage Employed</i>   | 8. <i>Earnings Children Health</i> |
| 3. <i>Advantage Employed Health</i>  | 9. <i>Employed Children Health</i> |
| 4. <i>Advantage Employed Family</i>  | 10. <i>Employed Health Family</i>  |
| 5. <i>Earnings Employed Children</i> | 11. <i>Employed Health Age</i>     |
| 6. <i>Earnings Employed Health</i>   |                                    |

Among these, the expert chose the fifth partition as he thought it constituted an important and comprehensible intermediate concept representing applicants' *Status*.

The decomposition process continued in a similar way resulting in three more aggregate attributes: *Social* (social and health condition of the applicant), *Present* (suitability of applicant's present housing) and *House* (overall housing conditions). The overall model structure is presented in Fig. 5. Apart from the two excluded attributes, the resulting structure is very similar to the one actually used in the Fund's decision support system. Thus, when evaluated by the domain expert, the reconstruction was considered successful.



**Fig. 5.** Model structure developed in supervised mode

The classification accuracy of the obtained model is very high:  $97.8 \pm 1.8\%$ . In comparison with the unsupervised decomposition, this is about 3% better, so the model developed in the supervised mode is superior. In other words, the model developed by a combination of DM and DS techniques was clearly better than the one developed by DM alone.

#### 6.4 Serial Mode

Serial mode of combining DM and DS occurs when HINT and DEX are used one after another contributing to a single model. First, an initial model is developed by HINT from data, and then modified by the expert using DEX. This is possible because the two systems use exactly the same model representation. Once a model has been developed by HINT, it is already represented in the form suitable for DEX.

For example, recall that two basic attributes, *Cult-Hist* and *Fin-Sources*, were excluded from the models developed by HINT. These attributes were not truly redundant; they were eliminated because they corresponded to some rare cases that had not occurred in the data set. In order to reintroduce them into the model, the expert can take a model developed by HINT (either in supervised or unsupervised mode), load it into DEX, manually add the two attributes, and edit the corresponding decision rules.

#### 6.5 Parallel Mode

Parallel mode refers to the case when DM and DS are used independently to make two or more different models, either from data by HINT, or from expertise by DEX. The result is a multitude of models, which can be compared with each other in terms of used attributes, their structure and inter-relations, or classification accuracy. In addition to a better insight into the problem, such models can be used to provide “second opinion”, which can be either confirmative when the evaluations agree, or can raise a warning in the case of disagreement.

Looking back along this paper, it is easy to see that we actually did use the parallel mode: we developed a model by DEX (Fig. 2) and two models by HINT (Figs. 4 and 5). In sections 6.2 and 6.3, we compared them in terms of their structure, comprehensibility, and classification accuracy.

#### 6.6 Combined Mode

Combined mode is the most general and subsumes some of the others. The idea is to develop a single model using sub-models (subtrees) possibly developed by different methods and from different sources. So, one or more subtrees can be developed by HINT, some others by DEX, and all combined together into a single model. Notice

that multi-attribute models are particularly suitable for this kind of development due to their hierarchical structure and inherent flexibility.

For a hypothetical example, consider the model from Fig. 2, which was originally in its entirety developed by DEX. The model contains three subtrees: *House*, *Status*, and *Soc-Health*. Had the appropriate data about earnings, employment, and social and health status of citizens been available prior to that, the subtrees *Status* and *Soc-Health* could have been obtained by HINT, possibly from two or more different data sets. The third subtree, *House*, could have been developed by, say, a real estate expert. After obtaining these components, they could have been manually "glued" together by introducing a new root attribute *Housing* and defining the corresponding decision rules.

## 7 Further Work

The loan allocation model presented in this paper was originally developed only by a DS method; the integration of DM and DS for the same problem was revisited afterwards in a somewhat hypothetical setting. Obviously, the approach needs further evaluation in real-life settings so as to combine DM and DS right from the beginning.

Apart from multi-attribute models, are there any other types of models suitable for the proposed principle of DM or DS integration? The approach requires that (1) there exist both DM and DS methods of model development that (2) use the same model representation. Probably most promising for further research seem representations based on production (if-then) rules. There has been a long tradition of developing such rules manually for expert systems [13], and there exist machine learning methods that can learn rules from data, including C4.5 [17]. Some potential problems may occur with the supervised mode of operation, which may not be possible with some rule-learning algorithms, and with the combined mode since production rules are in general not as well structured as hierarchical models.

Another interesting candidate for further research may be decision trees, which are widely used both in decision analysis (DA) and machine learning (ML). However, precaution is needed because these two disciplines use the same name for two somewhat different concepts. In ML, a decision tree consists of internal nodes that represent some branching criteria, and terminal nodes that correspond to the final classification or evaluation. In contrast, DA's decision trees contain chance and decision nodes, and their terminal nodes represent utilities, which are always numeric and typically represent some monetary value. In spite of the differences, there already exist DS systems that can deal with both types of decision trees, for example DATA [19], a computer program that facilitates the manual development of both DA and ML type decision trees. A corresponding machine learning algorithm is still missing, though.

Finally, the concepts that are manually developed in DEX can be used as background knowledge for some DM method. In general, DEX does not need to be coupled with HINT, as it can provide concepts for other types of learning methods. For instance, in [24] the concepts manually developed by a physician and encoded in DEX were used to complement the data when inducing a naive Bayesian prognostic

model. This model was significantly more accurate than the model developed from data only.

## 8 Conclusion

We proposed an approach that integrates Data Mining (DM) and Decision Support (DS) for model-based problem solving. In general, the approach requires to be capable of developing models in both DS and DM ways: “manually” from expertise, and “automatically” from data, respectively. It is also important that both types of methods develop models of the common type and use the same representation; this facilitates a multitude of method combinations, such as supervised, parallel, and serial. To attempt such a combination, both human expertise and mining data about the problem must be available.

Specifically, our approach is based on qualitative hierarchical multi-attribute decision models and two model development tools: DEX on the DS side, and HINT on the DM side. Using a real-world example of housing loan allocation, we demonstrated that the approach was indeed practical and offered a number of improvements of the modeling process. Some highlights include:

1. HINT successfully reconstructed the DEX’s model and outperformed C4.5 in terms of classification accuracy.
2. The integration of DM and DS through HINT’s supervised decomposition had a positive effect on both the classification accuracy and comprehensibility.
3. The parallel mode of operation provides a multitude of models and evaluations, including a so-called “second opinion”.
4. Hierarchical multi-attribute models are highly modular and flexible, and thus convenient for the combined mode of operation.

Regarding further work, this approach needs additional evaluation in real-world applications. Also, it seems promising for the exploration of other modeling formalisms, especially production rules and decision trees.

## Acknowledgment

The work reported here was supported by the Slovenian Ministry of Education, Science and Sport, and by the EU project SolEuNet, IST-11495.

## References

1. Andrienko, G., Jankowski, P., Andrienko, N.: Data mining approach to reducing complexity of multiple criteria spatial decision problem. In: Forer, P., Yeh, A.G.O., He, J. (eds.): Proc. 9th Int. Symp. on Spatial Data Handling SDH-2000. Beijing (2000) 3a15–3a28
2. Ashenurst, R.L.: The Decomposition of Switching Functions. Technical report, Bell Laboratories BL-1 (11) (1952) 541–602

3. Angehm, A.A.: Supporting multi-criteria decision making. In: Holtham C. (ed.): *Executive Information Systems and Decision Support*. Chapman & Hall (1992)
4. Bohanec, M., Rajkovič, V.: DEX: An expert system shell for decision support. *Sistemica* 1(1) (1990) 145–157
5. Bohanec, M., Cestnik, B., Rajkovič, V.: A management decision support system for allocating housing loans. In: Humpreys, P., Bannon, L., McCosh, A., Migliarese, P. (eds.): *Implementing Systems for Supporting Management Decisions*. Chapman & Hall (1996) 34–43
6. Bohanec, M., Cestnik, B., Rajkovič, V.: Evaluation models for housing loan allocation in the context of floats. In: Berkeley, D., Widmeyer, G.R., Brézillon, P., Rajkovič, V. (eds.): *Context Sensitive Decision Support Systems*. Chapman & Hall (1998) 174–189
7. Bohanec, M., Rajkovič, V.: Multi-attribute decision modeling: Industrial applications of DEX. *Informatica* 23 (1999) 487–491
8. Bohanec, M., Zupan, B., Rajkovič, V.: Applications of qualitative multi-attribute decision models in health care. *International Journal of Medical Informatics* 58–59 (2000) 191–205
9. Bohanec, M., Rajkovič, V., Leskošek, B., Kapus, V.: Expert knowledge management for sports talent identification and advising process. In: Carlsson, S.A., Brézillon, P., Humpreys, P., Lundberg, B.G., McCosh, A.M., Rajkovič, V. (eds.): *Decision Support through Knowledge Management* (2000) 46–59
10. Clemen, R.T.: *Making Hard Decisions: An Introduction to Decision Analysis*. Duxbury Press (1996)
11. Curtis, H.A.: *A New Approach to the Design of Switching Functions*. Van Nostrand (1962)
12. Han, J., Kamber, M.: *Data Mining: Concepts and Techniques*. Morgan Kaufman (2001)
13. Mallach, E.G.: *Understanding Decision Support Systems and Expert Systems*. Irwin (1994)
14. Mallach, E.G.: *Decision Support and Data Warehouse Systems*. McGraw-Hill (2000)
15. Microsoft Corporation: *OLE DB for Data Mining Specification*. <http://download.microsoft.com/download/dasdk/Install/1/WIN98Me/EN-US/oledbdm.exe> (2000).
16. Provost, F., Fawcett, T.: Robust classification for imprecise environments. *Machine Learning* 42 (2001) 203–231
17. Quinlan, J.R.: *C4.5: Programs for Machine Learning*. Morgan Kaufmann (1993)
18. Saaty, T.L.: *Multicriteria Decision Making: The Analytic Hierarchy Process*. RWS Publications (1993)
19. TreeAge Software, Inc.: *DATA 3.5 for Healthcare User's Manual*. TreeAge Software, Inc., Williamstown, MA, USA. <http://www.treeage.com> (1999)
20. Urbančič, T., Križman, V., Kononenko, I.: *Review of AI Applications*. Jožef Stefan Institute, Ljubljana, Report DP-7806 (1998)
21. Zupan, B., Bohanec, M., Demšar, J., Bratko, I.: Learning by discovering concept hierarchies. *Artificial Intelligence* 109 (1999) 211–242
22. Zupan, B., Bratko, I., Bohanec, M., Demšar, J.: Induction of concept hierarchies from noisy data. In: Langley, P. (ed.): *Proc. Int. Conf. on Machine Learning ICML-2000*. Morgan Kaufman (2000) 1199–1206
23. Zupan, B., Demšar, J.: *Orange*. <http://magix.fri.uni-lj.si/orange/index.htm> (2001)
24. Zupan, B., Demšar, J., Smrke, D., Božikov, K., Stankovski, V., Bratko, I., Beck, J.R.: Predicting patient's long term clinical status after hip arthroplasty using hierarchical decision modeling and data mining. *Methods of Information in Medicine* 40 (2001) 25–31