

Explaining clinical decisions by extracting regularity patterns

Concha Bielza^a, Juan A. Fernández del Pozo^{a,*}, Peter J.F. Lucas^b

^a *Decision Analysis Group, Technical University of Madrid, Campus de Montegancedo, Boadilla del Monte, 28660 Madrid, Spain*

^b *Institute for Computing and Information Sciences, Radboud University Nijmegen, Toernooiveld 1, 6525 ED Nijmegen, The Netherlands*

Received 30 September 2005; received in revised form 13 October 2006; accepted 6 May 2007

Available online 13 May 2007

Abstract

When solving clinical decision-making problems with modern graphical decision-theoretic models such as influence diagrams, we obtain decision tables with optimal decision alternatives describing the best course of action for a given patient or group of patients. For real-life clinical problems, these tables are often extremely large. This is an obstacle to understand their content. *KBM2L* lists are structures that minimize memory storage requirements for these tables, and, at the same time, improve their knowledge organization. The resulting improved knowledge organization can be interpreted as explanations of the decision-table content. In this paper, we explore the use of *KBM2L* lists in analyzing and explaining optimal treatment selection in patients with non-Hodgkin lymphoma of the stomach using an expert-designed influence diagram as an experimental vehicle. The selection of the appropriate treatment for non-Hodgkin lymphoma of the stomach is, as for many other types of cancer, difficult, mainly because of the uncertainties involved in the decision-making process. In this paper we look at an expert-designed clinical influence diagram as a representation of a body of clinical knowledge. This diagram can be analyzed and explained using *KBM2L* lists. It is shown that the resulting lists provide high-level explanations of optimal treatments for the disease. These explanations are useful for finding relationships between groups of variables and treatments. It is demonstrated that these lists can act as a basis for gaining a deeper understanding of the underlying clinical problem.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Medical decision-support systems; Decision analysis; Influence diagrams; Knowledge analysis; Knowledge explanation

1. Introduction

An *influence diagram* is a modern decision-theoretic formalism. Nowadays it is frequently adopted as a basis for constructing decision-support systems (DSS) and used to structure and solve decision-making problems [18]. It consists of an acyclic directed graph with associated probabilities and utilities modeling the

uncertainties and preferences tied in with the problem concerned. The result of solving (or evaluating as is the technical term) an influence diagram are *decision tables* containing the optimal decision alternatives. Thus, for every decision, there is an associated decision table with the best alternative, i.e. the alternative with the maximum expected utility for every combination of relevant variables that are observable before the decision is made. The evaluation algorithm determines which of the observable variables are relevant.

For some medical problems, usually problems that involve difficult trade-offs between the benefits and risks of a treatment, doctors may use decision tables to

* Corresponding author. Tel.: +34 91 3366596; fax: +34 91 3524819.

E-mail addresses: mcbielza@fi.upm.es (C. Bielza), jafernandez@fi.upm.es (J.A. Fernández del Pozo), peterl@cs.ru.nl (P.J.F. Lucas).

determine the best patient treatment recommendations. However, medical doctors may find it difficult to accept such recommendations if they do not understand the reasons behind them [9,12]. The questions a medical doctor is likely to ask are: “Why is the proposed decision optimal?” and “What are the implicit rules underlying the modeled decision problem?” Answering these questions can be seen as providing *explanations* to medical decisions. These answers may provide new insights into the problem, and, as a form of knowledge synthesis, may also be useful for validating a system.

Considering that the table sizes are exponential in terms of the number of variables, the search for explanations is not an easy task from a purely computational viewpoint either. Turning the huge tables into more compact formats will lead to memory savings. Although finding explanations and optimizing the storage space of the decision tables are obviously not the same problem, both challenges can be addressed simultaneously as the sought-after compactness implies searching for those explanations.

In [7], we introduced *KBM2L* lists to address this problem. Finding explanations is a goal pursued within a number of disciplines, such as knowledge-based systems and machine learning. Thus, our approach bears some resemblance to knowledge extraction techniques, such as are used to construct tree-based classifiers [3], oblivious read-once decision graphs [11], rough sets [17], and to identify which nodes are relevant for each decision node in an influence diagram [13,20]. As explained in detail in [7], the *KBM2L* method tries to reorganize a knowledge structure by a global search for

good, representative candidates. At the start of the algorithm there are already correctly classified cases represented in table form. These cases can be interpreted as representing types of patients, and we try to extract the reasons underlying this classification. Unlike the typical situation in machine learning, these cases are unique as they correspond to configurations of variables in the original influence diagram. Note that our method is applied after the influence diagram has been solved, whereas most of the work on the influence diagram framework concerns operations on the graph structure before evaluating the influence diagram [13].

To investigate the practical usefulness of this method within a medical setting, an influence diagram regarding the treatment of non-Hodgkin lymphoma of the stomach was chosen as an experimental vehicle [14]. This is a realistic clinical model, reflecting the current scientific evidence from the medical literature about this disorder. It can be used to determine the optimal treatment for individual patients with non-Hodgkin lymphoma of the stomach. A treatment involves deciding whether or not to prescribe antibiotics, whether to undertake curative or palliative surgery or no surgery at all, and, finally, which combination of chemotherapy, radiotherapy, if any. The probabilistic part of the model, a Bayesian network, can be used independently to predict prognosis and to generate patient-specific risk profiles. The model exceeds common prognostic models based on logistic regression, as it is part of a DSS that can answer many different clinical questions [15]. In this paper, we analyze this influence diagram with the aim of extracting patterns. These patterns are then used to explain the

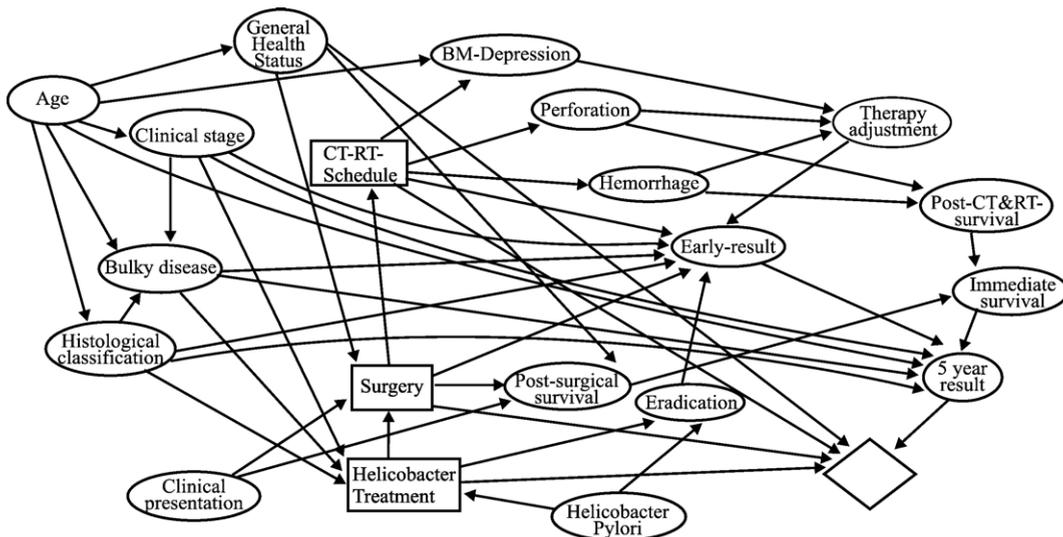


Fig. 1. Influence diagram for the treatment of gastric NHL.

optimal treatment alternatives that can be generated by evaluating the diagram.

The paper is organized as follows. In Section 2, we review the clinical problem of treating gastric non-Hodgkin lymphoma and in Section 3 we discuss the technique of *KBM2L* lists and its role in the analysis of the gastric non-Hodgkin lymphoma treatment model. Next, Section 4 describes the results of applying *KBM2L* lists to the gastric non-Hodgkin lymphoma model. Finally, the paper is rounded off in Section 5 with some conclusions.

2. Treatment selection for patients with gastric non-Hodgkin lymphoma

Primary gastric non-Hodgkin lymphoma, *gastric NHL* for short, is a relatively rare disorder, accounting for about 5% of gastric tumors. This disorder is caused by a chronic infection by the *Helicobacter pylori* bacterium, *H. pylori* for short [4]. Treatment consists of a combination of antibiotics, chemotherapy, radiotherapy and surgery. The details are discussed below. In particular for patients with a poor general health status and with advanced disease, it is difficult to select an appropriate treatment. Many clinicians, therefore, believe that some form of decision support is needed to assist clinicians with the treatment of these patients. Models that have been used for this purpose in the past are normally only capable of predicting the prognosis of the disease, but cannot be used to select treatment.

To overcome these limitations, the last author with the help of two oncologists constructed a number of influence diagrams [14]. These models are only meant to be used for patients with histologically confirmed gastric NHL. We have taken the most complex version with 3 decision nodes for this study. This influence diagram is shown in Fig. 1, and is briefly discussed in the following. The first of the decision nodes, helicobacter-treatment (HT), corresponds to the decision to prescribe antibiotics against *H. pylori*. The second decision concerns carrying out surgery (S). The possibilities are either curative surgery, involving the complete removal of the stomach and locoregional tumor mass, palliative surgery, i.e. partial removal of the stomach and tumor, or no surgery. The last decision, ct–rt-schedule (CTRTS), is concerned with the selection of chemotherapy (Chemo), radiotherapy (Radio), chemotherapy followed by radiotherapy (Ch.Next.Rad), or none.

The influence-diagram model consists of 17 chance nodes (ellipses), one value node (diamond), three decision nodes (rectangles), 42 arcs, 8282 probability entries where the largest table has 3840 entries, and an initial table for the value node of size 144. Nodes to the

left of the decision nodes (see Fig. 1) concern pre-treatment information. Nodes to the right of the decision nodes are posttreatment nodes. Some of the variables with their associated domain are listed in Table 1.

The diagram structure was determined from the known causal relationships and the probabilistic (in) dependencies found in the literature. Probabilities were elicited with the aid of logical and qualitative probabilistic relationships. The numerical assessments were checked against these relationships. For the two largest conditional probability tables associated with early-result and 5-year-result, a special additive model was used to ease the assessment.

A Bayesian network was obtained from the influence diagram by converting the decision nodes into chance nodes [2]. Thus, the accuracy of the assessments could be tested by comparing prior and posterior marginal probability distributions of this network with frequency data taken from the literature and clinical experience. Also, a database with 137 patients was employed to

Table 1
The gastric NHL variables with their possible values

Variable	Possible values
Helicobacter-treatment (HT)	No, Yes
Surgery (S)	None, Curative, Palliative
ct–rt-schedule (CTRTS)	None, Radio, Chemo, Ch.Next.Rad
General health status (GHS)	Poor, Average, Good
Clinical stage (CS)	I, II1, II2, III, IV
bulky-disease (BD)	Yes, No
Histological-classification (HC)	Low-grade, High-grade
<i>Helicobacter pylori</i> (HP)	Absent, Present
Clinical-presentation (CP)	None, Hemorrhage, Perforation, Obstruction
Age	v10.19, v20.29, v30.39, v40.44, v45.49, v50.54, v55.59, v60.64, v65.69, v70.79, v80.89, GE90
Eradication	No, Yes
Bm-depression (Bone marrow)	No, Yes
Perforation	No, Yes
Hemorrhage	No, Yes
Therapy-adjustment	No, Yes
Post-ct–rt-survival	No, Yes
Post-surgical-survival	No, Yes
Immediate-survival	No, Yes
Early-result	CR (complete remission), PR (partial remission), NC (no change), PD (progressive disease)
5-year-result	ALIVE, DEATH

assess the model's accuracy. Finally, a new joint probability distribution was learned from the database, and its prior marginal probabilities were compared with those of the original network.

Two methods were used to elicit utilities from the patient's perspective: direct scaling and the reference gamble. In the first method, the possible clinical outcomes are assessed directly using a linear numerical scale. The second method is based on gambles or choices between lotteries making the assessment indirect. Standard reference texts describe these methods [5]. Several utility functions were obtained and refined by examining the performance of the system with respect to the optimal treatments proposed for some typical patients. This yielded a clinically reliable utility function.

A preliminary evaluation of the gastric NHL model's accuracy has already been done by means of a double blind clinical study. This research, where we use *KBM2L* lists to get a better understanding of the treatment basis of the gastric NHL model, adds to this earlier study. The encouraging results we obtained for another medical problem inspired us to investigate the gastric NHL diagram [6].

3. Explanation of decision options by *KBM2L* lists

Next we look at how *KBM2L* lists work. *KBM2L* lists are investigated in this paper with regard to their explanatory power.

3.1. Basics

A decision-table output by evaluating an influence diagram comprises two parts: (1) a set of all variable configurations, which can be indexed assuming a natural or conventional order, in the values of their discrete domains; (2) the table content, i.e. the optimal alternative. As variables in the clinical field are usually patient attributes, the term *attributes* will be used in the following instead of 'variables'.

The attributes can be arranged in different orders, maintaining however the same information. A *base* is defined as a vector with elements equal to the attributes in a specific order. Given a base, an *index* is a vector whose elements are the base attribute values, interpreted as the coordinates with respect to that base. If we have a fixed order of attributes with discrete domains, we can view a decision table as a multidimensional matrix.

We can map this multidimensional matrix to a linear array or list in a similar way to sequential memory allocation in computers [10]. Given a cell of the table

with index $c=(c_0, c_1, \dots, c_n)$, we define the access function $f: R^{n+1} \rightarrow R$, such that

$$f(c_0, c_1, \dots, c_n) = c_0 \prod_{i=1}^n D_i + c_1 \prod_{i=2}^n D_i + \dots + c_n = q \quad (1)$$

where q is the c -offset with respect to the first element of the table in a given base, and D_i denotes the cardinality of the i th attribute domain for $i=0, 1, \dots, n$. The access function f can also be written more compactly as

$$f(c_0, c_1, \dots, c_n) = \sum_{i=0}^n c_i w_i \quad (2)$$

where $w_i = \prod_{j=i+1}^n D_j = w_{i+1} D_{i+1}$ is called the weight of the i th attribute, $i=0, 1, \dots, n$.

Thus, index-notation and offset-notation are equivalent, and are related to each other by the function f defined by Eqs. (1) or (2).

To shorten the list output by the decision table, we look at the cell content, i.e. the optimal alternatives. It is not unusual to find that some *consecutive* cells lead to the *same* optimal alternative. The number of such consecutive cells represents the knowledge granularity of optimal decisions. Then, a new compact list can be constructed, much in the same way as sparse matrices are managed. This list will only store one index (or equivalently its offset) per set of equal alternatives. We will choose the last index (offset) as the representative of this set. This last index together with the shared optimal alternative, representing a set of cases, is called an *item*. The resulting list of items is called a *KBM2L* list, which stands for a "Knowledge Base Matrix to List" representation [7].

An item is denoted by $\langle \text{index}, \text{alternative} \rangle$ or, equivalently, $\langle \text{offset}, \text{alternative} \rangle$, where ' \langle ' reflects that the item offsets increase monotonously, and ' \rangle ' reflects granularity. For example, the length of the fragment of the list in offset notation $\langle p-1, y \rangle \langle p, x \rangle \dots \langle p+q, x \rangle \langle p+q+1, z \rangle$, with x, y, z three different optimal alternatives, is $q+3$. This fragment can be collapsed to $\langle p-1, y \rangle \langle p+q, x \rangle \langle p+q+1, z \rangle$ as a *KBM2L* list with length 3.

Consider a set of indices representing an item. Since the indices are ordered, this set will range from an index I_{inf} to I_{sup} , corresponding to the extreme cases (infimum and supremum) of the item. In this set of indices we can consider a *fixed* part, representing the index components common to all the item cases, and a *variable* part, where the values of the attributes corresponding to the indices are not shared. Both parts can be derived from the indices I_{inf} and I_{sup} , e.g., the fixed part is obtained by taking the logical AND: $I_{\text{inf}} \vee I_{\text{sup}}$.

These concepts open up ways to automatically generate explanations for decisions. The fact that the values of the attributes in the fixed part of items are equal somehow *explains* why the optimal alternative is also the same across items. Hence, the set of attributes of the *fixed part can be interpreted as explaining the reasons behind the optimal alternative*. The attributes in the variable part of an item are irrelevant for decision-making.

Example 1. This simple example illustrates the basics. Suppose that the decision table after evaluating an influence diagram is as follows:

X_1	X_2	Offset	Policy
0	0	0	B
0	1	1	A
1	0	2	B
1	1	3	A

That is, we have two binary attributes X_1 and X_2 , and $\{A, B\}$ is the set of alternatives. There exist two possible bases: $[1,2]$ and $[2,1]$. The first row of the table has the $(0, 0)$ index for both bases. The second row of the table has the $(0, 1)$ index for base $[1,2]$ and the $(1, 0)$ index for base $[2,1]$, and so on.

Given the $[1,2]$ base, the access function becomes $f(c_1, c_2) = c_1 \cdot 2 + c_2 = q$. This formula has been used to compute the *offset* column in the table. This table is then mapped to list $\langle\langle(0, 0), B\rangle\rangle \langle\langle(0, 1), A\rangle\rangle \langle\langle(1, 0), B\rangle\rangle \langle\langle(1, 1), A\rangle\rangle$, expressed in index-notation. Its offset-notation is $\langle(0, B)\rangle \langle(1, A)\rangle \langle(2, B)\rangle \langle(3, A)\rangle$. The derived *KBM2L* list has 4 items: $\langle 0, B | \langle 1, A | \langle 2, B | \langle 3, A |$ (in offset-notation).

However, given the $[2,1]$ base, with access function $f(c_2, c_1) = c_2 \cdot 2 + c_1$, the linear array is $\langle\langle(0, 0), B\rangle\rangle \langle\langle(0, 1), B\rangle\rangle \langle\langle(1, 0), A\rangle\rangle \langle\langle(1, 1), A\rangle\rangle$ in index-notation, or $\langle(0, B)\rangle \langle(1, B)\rangle \langle(2, A)\rangle \langle(3, A)\rangle$ in offset-notation. It results in a *KBM2L* list with only 2 items: $\langle 1, B | \langle 3, A |$ (in offset-notation). The fixed part of (c_2, c_1) for the first item $\langle 1, A |$ is c_2 since $I_{\text{inf}} = (c_2 = 0, c_1 = 0)$ and $I_{\text{sup}} = (c_2 = 0, c_1 = 1)$. c_1 is the variable part. The item contains 2 cases and $c_2 = 0$ explains the policy B .

3.2. Implementation

We have implemented the process of building a *KBM2L* list from a decision table as outlined above. The process starts with an empty list, i.e. with an item representing the complete absence of knowledge. Given a base, each case is sequentially added to the list, which is implemented employing up to 26 rules for item management. Each rule examines what the *KBM2L* list of items is like before and after applying the rule.

Different bases may contain the same table knowledge. However, the granularity, data organization, and, possibly, the memory requirements to store the final list of items may vary from one base to another. Our aim is to get a base that minimizes the number of items, bringing up the granular knowledge. Searching for solutions in the space of possible attribute permutations with a fixed domain order is a combinatorial optimization problem, which is known to be NP-hard [8]. During the search, it may be very time-consuming, or even intractable to compute the new storage space that each base requires and copy the information from one *KBM2L* to another in a different base. In [7], we discuss efficient heuristics to deal with such problems. To drive the search for a good base, a genetic and a variable neighborhood algorithm have been implemented and tested. In addition, there are a number of learning heuristics for reorganizing and rapidly, possibly partially, copying complex lists. Furthermore, the search space may be reduced by discarding candidates for better bases. This is done by means of a procedure that infers whether a given list is inferior to the present one.

Sometimes the decision table is too large to be fully evaluated. In this case, a set of subproblems is solved instead. Each subproblem is the result of instantiating some random variables. This subproblem set may not be exhaustive. This implies that there will be unknown optimal alternatives for some attribute combinations, i.e. for combinations associated with unsolved subproblems. This is not a problem for the *KBM2L* construction process which also operates with unknown policies. First, it evaluates, sequentially or in parallel, all the subproblems. Then, the resulting partial decision tables are sequentially added to the *KBM2L* list by means of a learning mechanism that optimizes the list before processing the next partial table. Each stage in the addition process improves the item organization and facilitates future additions [7].

Example 2. The ideas summarized in this subsection are illustrated by means of a simple decision problem stored in two tables. The first table is an extension of the table given in Example 1. There are three different attributes X_0, X_1 and X_2 . It is assumed that the base is equal to $[0, 1, 2]$, and that all the domains are binary, with possible values 0 and 1. The set of alternatives is equal to $\{A, B, C\}$. Table 2 represents the evaluation output for the instance with $X_0 = 0$. Table 3 lists the results for $X_0 = 1$.

The initial empty list is equal to $\langle 7, -1 |$, where -1 means no knowledge, and 7 is the result of having to consider 8 cases, where counting starts at 0. Firstly,

Table 2
Partial decision table ($X_0=0$)

X_0	X_1	X_2	Offset	Policy
0	0	0	0	<i>B</i>
0	1	0	2	<i>B</i>
0	1	1	3	<i>A</i>
0	0	1	1	<i>A</i>

Table 2 is translated to the *KBM2L*. We consider each case in the table and add them one by one into the *KBM2L*.

The sequence of steps is as follows:

- (1) Case $((0, 0, 0), B)$ at offset 0 yields $\langle 0, B \mid \langle 7, -1 \rangle$;
- (2) Case $((0, 1, 0), B)$ at offset 2 yields $\langle 0, B \mid \langle 1, -1 \rangle \langle 2, B \mid \langle 7, -1 \rangle$ (where $\langle 1, -1 \rangle$ acts as an empty placeholder);
- (3) Case $((0, 1, 1), A)$ at offset 3 yields $\langle 0, B \mid \langle 1, -1 \rangle \langle 2, B \mid \langle 3, A \mid \langle 7, -1 \rangle$;
- (4) Case $((0, 0, 1), A)$ at offset 1 yields $\langle 0, B \mid \langle 1, A \mid \langle 2, B \mid \langle 3, A \mid \langle 7, -1 \rangle$.

This list is fragmented, as it consists of 5 items. Thus, there is room for optimization. We call $\langle 0, B \mid \langle 1, A \mid \langle 2, B \mid \langle 3, A \mid \langle 7, -1 \rangle$ the *initial* list. Because this is a small example, we can simply test all the possible solutions, i.e. $3! = 6$ solutions. The best base is $[0, 2, 1]$, with an associated list of 3 items: $\langle 1, B \mid \langle 3, A \mid \langle 7, -1 \rangle_{[0,2,1]}$. For the sake of clarity, the base is shown as a subscript of the item list.

Subsequently, the other partial results (cf. Table 3) are added to the optimized list of three items in the following order:

- (5) Case $((1, 0, 1), C)$ at offsets $5_{[0,1,2]}$ and $6_{[0,2,1]}$ yielding $\langle 1, B \mid \langle 3, A \mid \langle 5, -1 \rangle \langle 6, C \mid \langle 7, -1 \rangle_{[0,2,1]}$;
- (6) Case $((1, 1, 0), C)$ at offsets $6_{[0,1,2]}$ and $5_{[0,2,1]}$ yielding $\langle 1, B \mid \langle 3, A \mid \langle 4, -1 \rangle \langle 6, C \mid \langle 7, -1 \rangle_{[0,2,1]}$;
- (7) Case $((1, 0, 0), A)$ at offsets $4_{[0,1,2]}$ and $4_{[0,2,1]}$ yielding $\langle 1, B \mid \langle 4, A \mid \langle 6, C \mid \langle 7, -1 \rangle_{[0,2,1]}$;
- (8) Case $((1, 1, 1), C)$ at offsets $7_{[0,2,1]}$ and $7_{[0,2,1]}$ yielding $\langle 1, B \mid \langle 4, A \mid \langle 7, C \mid_{[0,2,1]}$.

The list $\langle 1, B \mid \langle 4, A \mid \langle 7, C \mid_{[0,2,1]}$ is now the final list and is subsequently optimized. Since there are three

Table 3
Partial decision table ($X_0=1$)

X_0	X_1	X_2	Offset	Policy
1	0	1	5	<i>C</i>
1	1	0	6	<i>C</i>
1	0	0	4	<i>A</i>
1	1	1	7	<i>C</i>

Table 4
Decision table after re-ordering

X_0	X_2	X_1	Offset	Offset	Policy
			$[0,1,2]$	$[0,2,1]$	
0	0	0	0	0	<i>B</i>
0	0	1	2	1	<i>B</i>
0	1	0	1	2	<i>A</i>
0	1	1	3	3	<i>A</i>
1	0	0	4	4	<i>A</i>
1	0	1	6	5	<i>C</i>
1	1	0	5	6	<i>C</i>
1	1	1	7	7	<i>C</i>

alternatives, and we have three items, the final list, $\langle 1, B \mid \langle 4, A \mid \langle 7, C \mid_{[0,2,1]}$, is the best one. Table 4 shows the associated decision table.

The so-called *KBM2L* spectrum chart is a useful aid for displaying the *KBM2L* optimization process. Such spectra depict the unidimensional memory layout and the base-dependent case grouping into items. The cases are ordered by the offset (X -axis), where cases with the same optimal alternative are shaded by the same color (or gray scale). The emerging role of data visualization has been identified as a separate and important data mining task [19]. Our spectrum falls into the class of pixel-oriented techniques. Fig. 2 shows this visualization tool for the initial item list $\langle 0, B \mid \langle 1, A \mid \langle 2, B \mid \langle 3, A \mid \langle 7, -1 \rangle_{[0,2,1]}$ (left-hand side of the figure) and the final item list $\langle 1, B \mid \langle 4, A \mid \langle 7, C \mid_{[0,2,1]}$ (right-hand side of the figure).

Finally, as an illustration of the two parts of an item, let us examine the last *KBM2L* list, consisting of 3 items. For item $\langle 1, B \mid$, it holds that $I_{\text{inf}} = (c_0=0, c_2=0, c_1=0)$ and $I_{\text{sup}} = (c_0=0, c_2=0, c_1=1)$. The item contains 2 cases. Consequently, (c_0, c_2) is the fixed part of the index, whereas c_1 is the variable part. The tuple $(c_0=0, c_2=0)$ to some extent explains the reasons underlying the policy *B*. For item $\langle 4, A \mid$, we have that $I_{\text{inf}} = (c_0=0, c_2=1, c_1=0)$ and $I_{\text{sup}} = (c_0=1, c_2=0, c_1=0)$. The item includes 3 cases. The fixed part is empty; the variable part is equal to the entire index. For item $\langle 7, C \mid$, $I_{\text{inf}} = (c_0=1, c_2=0, c_1=1)$ and $I_{\text{sup}} = (c_0=1, c_2=1, c_1=1)$. The item contains 3 cases and the fixed part is equal to c_0 .

3.3. Aim of the analysis of the gastric NHL influence diagram

Although it is complicated and time-consuming to develop a decision-theoretic model, such as the one discussed in Section 2 and also entails computational difficulties at the evaluation stage [1], such models have been shown to be very useful. They can be applied to yield prognostic information about a specific patient, given



Fig. 2. Initial (left) and final (right) KBM2L spectra. Every colored band represents an item.

particular patient characteristics, and to help to make therapeutic decisions. By examining all possible therapeutic decisions, taking into account their pros and cons, including patient preferences, it is also possible to determine which therapeutic decisions yield optimal results. Furthermore, by backward reasoning (assuming that the final results of the treatment are known), the models can be used to generate probabilistic profiles for

groups of patients that fit these final results. Obviously, all these are valuable capabilities for driving the clinical decision-making process. However, understanding the treatment advice generated by a DSS for the whole patient population is not so straightforward. Clinicians would benefit from having clear and concise explanations of the results output by the system. Such explanations would justify and improve the understanding of these results. In

HELICOBACTER-TREATMENT (HT)

$$B = [CS, BD, HC, HP]:$$

17 items



SURGERY (S)

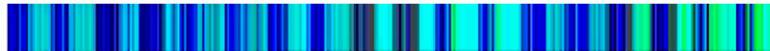
$$B = [GHS, HT, CS, BD, HC, HP, CP]:$$

385 items



CT-RT-SCHEDULE (CTRSTS)

$$B = [GHS, S, HT, CS, BD, HC, HP, CP]: 678 \text{ items}$$



HELICOBACTER-TREATMENT (HT)

$$B = [HP, HC, CS, BD]:$$

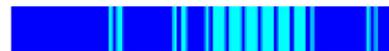
5 items



SURGERY (S)

$$B = [CP, HT, GHS, BD, CS, HC, HP]:$$

37 items



CT-RT-SCHEDULE (CTRSTS)

$$B = [CS, GHS, HC, BD, S, HT, CP, HP]: 369 \text{ items}$$



Fig. 3. Non-optimized (above) and optimized (below) KBM2L spectrum for each table.

Table 5
Combination of two decision tables

		HT	A	B	S					
		x0	a0	b0	y1 *					
		x0	a0	b1	y1 *	A	B	(HT,S)		
A	HT	x0	a1	b0	y1	a0	b0	(x1, 1)		
a0	x1	\oplus	x0	a1	b1	y2	=	a0	b1	(x1, 1)
a1	x0		x1	a0	b0	y1		a1	b0	(x0, 1)
			x1	a0	b1	y1		a1	b1	(x0, 2)
			x1	a1	b0	y2 *				
			x1	a1	b1	y2 *				

addition, this would be an alternative way for validating a system. In the remainder of the paper we investigate how *KBM2L* lists can be used for this purpose.

4. Results

In this section, we discuss the results of analyzing the gastric NHL influence diagram.

4.1. Decision tables

Evaluation of the influence diagram yielded three decision tables, one for each decision variable, each containing the optimal treatment for each combination of attributes in the tables. The first table, concerning the HT decision, depends on 4 attributes: CS, BD, HC and HP. We used the attributes in this order, i.e. base [CS, BD, HC, HP], to build a *KBM2L* list. It consisted of 17 items covering the whole set of 40 cases: 32 cases for HT=No, divided into 9 items, and 8 cases for HT=Yes, grouped into 8 items.

The second table, for the S decision, consisted of 7 attributes. The previous decision HT and two new attributes, GHS and CP, were added to this attribute set. The base [GHS, HT, CS, BD, HC, HP, CP] led to a *KBM2L* list with 385 items, covering the whole set of 960 cases. The distribution of these items was: 193 were associated with S=None (663 cases), and 192 with S=Curative (297 cases).

The last table concerning the CTRTS decision had 8 attributes. The *KBM2L* list built for the base [GHS, S, HT, CS, BD, HC, HP, CP] consisted of 678 items covering the whole set of 2880 cases. There were 164 items with CTRTS=None (490 cases), 188 with CTRTS=Radio (862 cases), 280 with CTRTS=Chemo (1404 cases) and 46 with CTRTS=Ch.Next.Rad (124 cases).

The three associated spectra are displayed at the top of Fig. 3. Every colored band in every spectrum represents an item. This illustrates the sensitivity of the data to changes in the context of a decision, and we can thus explore patterns.

Next, we applied our variable neighborhood algorithm to each *KBM2L* list to improve the bases. Improved lists for each decision node are shown at the bottom of Fig. 3. For HT, S, and CTRTS, the CPU times were 14.8, 144.7, and 2453.1 s, respectively, and 30, 124, and 223 base changes were required.

The three tables associated with the above lists were combined successively to produce a single, global table summarizing all the knowledge contained in the influence diagram. By combining the HT table (the first decision) with the S table (the second decision), the HT table is incorporated into the S table, which already includes HT as an attribute. The S table includes all the possible cases, even some cases known to be non-optimal from the ht decision table. These cases are therefore marked. Subsequently, a new table is constructed and represented as (HT, S). This table includes all the non-marked cases associated with the s table attributes. This is illustrated in Table 5. Here *A* and *B* are hypothetical binary attributes, \oplus denotes the operator that combines two tables, and the symbol * is used to mark a non-optimal case.

The S table was embedded into the CTRTS table similarly. The rows that did not match either the s table or the non-optimal rows were marked. All the marked rows were removed from the resulting table. The final, resulting table included 6 attributes, as (HT, S, CTRTS) formed the new combined optimal alternatives.

The resulting base was equal to $B_0 = [GHS, CS, BD, HC, HP, CP]$, with an associated *KBM2L* list consisting of 320 items. Note that this global table potentially contains more possible decisions than the simpler lists related to the individual treatments, i.e. up to $2 \cdot 3 \cdot 4 = 24$. However, only 14 were actually obtained, as the

Table 6
Optimizing the global table

Optimal alternative $T = (HT, S, CTRTS)$	# Cases	# Items in B_0	# in B_{final}	#T-items in B_T
(No, None, None)	45	28	30	24
(No, None, Radio)	72	48	37	5
(No, None, Chemo)	151	76	50	39
(No, Curative, None)	12	12	6	6
(No, Curative, Radio)	59	40	26	25
(No, Curative, Chemo)	39	39	31	10
(No, Curative, Ch.Next.Rd)	6	6	6	3
(Yes, None, None)	19	13	4	4
(Yes, None, Radio)	22	16	6	6
(Yes, None, Chemo)	21	11	12	4
(Yes, None, Ch.Next.Rad)	2	2	2	1
(Yes, Curative, None)	8	8	4	4
(Yes, Curative, Radio)	18	15	6	6
(Yes, Curative, Chemo)	6	6	4	4
Total	480	320	224	171

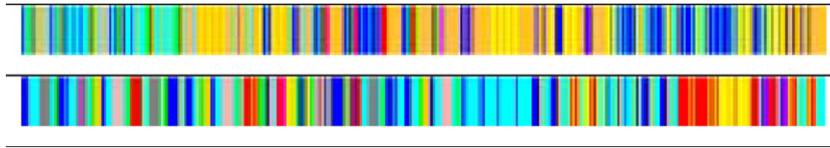


Fig. 4. Spectra for the global table: non-optimized spectrum (above), with B_0 =[GHS, CS, BD, HC, HP, CP] and 320 items; and optimized spectrum (below), with B_{final} =[HC, CP, GHS, HP, CS, BD] and 224 items.

remaining rows did not match. Palliative surgery was never selected, either separately or in combination with other treatments. This indicates that there is a bias in the diagram towards treating patients to cure: even patients with a very poor prognosis, who might have benefited from palliative surgery, were given a curative treatment. The information of this *KBM2L* is shown in the first three columns of Table 6.

Taking into account the cardinality of each attribute domain, the number of possible combinations was 480. Therefore, the knowledge represented by the 320 items appeared to be considerably fragmented. There were therefore reasons to consider optimizing the table. After 96 base changes run in 593.9 s, we obtained a shorter list, which refined the knowledge about the decisive attributes. The list consisted of 224 items, a significant improvement (a 30% reduction). The optimal base was equal to B_{final} =[HC, CP, GHS, HP, CS, BD]. The item distribution is shown in the fourth column of Table 6. The last column refers to B_T , a base obtained when the aim is just to explain a specific treatment T . This base will be explained in Section 4.3.

Fig. 4 shows the two spectra associated with B_0 and B_{final} .

4.2. Clinical interpretation of rules

Table 7 presents a portion of the optimal *KBM2L* list. Next we will consider some of the more noteworthy items. The fixed part of each item, i.e. its explanation, is shown in bold face. This list contains a selection of the 224 items, now interpreted as clinical rules, indicating the optimal global policy in the consequent as a function of the key attributes shown in the rule’s antecedent. The further to the left the attribute is, the more important it becomes (as it gains a higher weight with respect to the base according to Eq. (1)). We will then discuss some of the generated rules, and offer a clinical interpretation.

Consider rule 185 (HT=No, S=Curative, CTRTS=Chemo) compared to rule 191 (HT=No, S=Curative, CTRTS=Ch.Next.Rad). Both rules include only one case, since their variable parts are empty. The GHS patient variable values are ‘Average’ and ‘Good’, respectively,

whereas the HP variable values are ‘Present’ and ‘Absent’, respectively. It also appears that a patient who is in good shape (GHS is equal to ‘Good’) is expected to be able to cope with more intensive treatment, i.e. surgery and chemotherapy followed by radiotherapy, than a patient in poor condition, who will only receive surgery and chemotherapy. Even if the *H. pylori* bacterium is present, no antibiotics are prescribed in either case, as this would make no sense after removal of the stomach in which the bacterium resides.

Table 7
The optimal *KBM2L* list obtained with B_{final} =[HC, CP, GHS, HP, CS, BD]

# Item	Description	Size
0	<HC: Low-grade, CP: None, GHS: Poor, HP: Absent, CS: {I, II}, BD:—, (HT: No, S: None, CTRTS: Radio)>	4
1	<HC: Low-grade, CP: None, GHS: Poor, HP: Absent, CS: {II2, IV}, BD: —, (HT: No, S: None, CTRTS: Chemo)>	6
⋮	⋮	⋮
84	<HC: Low-grade, CP: Perforation, GHS: Good, HP: Present, CS: III, BD: No, (HT: No, S: Curative, CTRTS: Chemo)>	1
85	<HC: Low-grade, CP: Perforation, GHS: Good, HP: Present, CS: IV, BD: Yes, (HT: Yes, S: None, CTRTS: Chemo)>	1
86	<HC: Low-grade, CP: Perforation, GHS: Good, HP: Present, CS: IV, BD: No, (HT: Yes, S: Curative, CTRTS: Chemo)>	1
⋮	⋮	⋮
185	<HC: High-grade, CP: Perforation, GHS: Average, HP: Present, CS: IV, BD: No, (HT: No, S: Curative, CTRTS: Chemo)>	1
⋮	⋮	⋮
191	<HC: High-grade, CP: Perforation, GHS: Good, HP: Absent, CS: IV, BD: No, (HT: No, S: Curative, CTRTS: Ch.Next.Rad)>	1
⋮	⋮	⋮
222	<HC: High-grade, CP: Obstruction, GHS: Good, HP: Present, CS: III, BD: No, (HT: No, S: None, CTRTS: Radio)>	1
223	<HC: High-grade, CP: Obstruction, GHS: Good, HP: Present, CS: IV, BD:—, (HT: No, S: None, CTRTS: Chemo)>	2

The *KBM2L* item list summarizes the whole ordered table of optimal decisions based on the last case in every sequence of cases with the same treatment. The notation is: <attribute list: values, (decision vector)>. The explanation of each item is shown in bold face. The decision vector includes three components: HT, S and CTRTS, one for each decision.

The difference between rule 84 (HT=No, S=Curative, CTRTS=Chemo) and rule 86 (HT=Yes, S=Curative, CTRTS=Chemo) can be explained by noting that the clinical stage (CS) of the disease is different for both items. It has been decided that the treatment for the most advanced stage of the disease (CS=IV) for the slowly progressing low-grade version of gastric NHL should be more complete than for the less advanced stage (CS=III), where both disease stages are essentially incurable. The selection of curative surgery is based on the presence of perforation of the stomach, which leaves little choice but to remove the stomach. Whether the prescription of antibiotics against the *H. pylori* bacterium (HT=Yes) improves the patient's life expectancy at clinical stage IV (CS=IV) is likely to move oncologists to debate, as there is no scientific evidence available at the moment to support either the prescription or non-prescription of antibiotics to these patients.

Alternatively, consider item 223, shown in Table 7. It signifies that, according to the gastric NHL influence diagram, chemotherapy is the appropriate treatment for patients with high-grade tumors, with *H. pylori* and gastric obstruction. The trade-off in this case is between curative surgery, which has the advantage of removing the cause of the obstruction and the drawback of having many side-effects, or no surgery. Combining surgery with chemotherapy is too intensive for many patients. Antibiotics have a slightly positive effect at a very early stage of the disease, but their contribution to the overall effectiveness of the treatment is small. As chemotherapy is effective in patients with high-grade gastric NHL, this is the treatment of choice in this case.

4.3. Treatment selection

Clinicians may be interested not only in examining and comparing rules as discussed above to understand the reasons behind treatments, but also in finding out under which circumstances a treatment is chosen as being the best. This requires a study examining the conditions for the selection of treatments for individual patients. Consider, for example, treatment $T \equiv (\text{HT}=\text{No}, \text{S}=\text{Curative}, \text{CTRTS}=\text{Ch.Next.Rad})$, which belongs to 6 items, as shown in Table 6, including item 191 considered earlier. Our aim now is to explain when and why T is chosen as being the best. By focusing on this specific treatment T , the list only has to distinguish the possible response T from any other treatment $\neg T$ different from T . This binary response $KBM2L$ list may achieve shorter (and better) explanations because if two items corresponding with $\neg T$ ($\neg T$ -items) manage to join when the base changes, then another two T -items will also be joined. The new base will

Table 8

The optimal $KBM2L$ list for treatment $T \equiv (\text{HT}=\text{No}, \text{S}=\text{Curative}, \text{CTRTS}=\text{Ch.Next.Rad})$ obtained with $B_T = [\text{HC}, \text{CP}, \text{GHS}, \text{BD}, \text{CS}, \text{HP}]$

# Item	Description	Size
0	$\neg T$	142
1	$\langle \text{HC: Low-grade}, \text{CP: Perforation}, \text{GHS: Average}, \text{BD: Yes}, \text{CS: II1}, \text{HP: Absent}, (T) \rangle$	1
2	$\neg T$	142
3	$\langle \text{HC: Low-grade}, \text{CP: Perforation}, \text{GHS: Average}, \text{BD: Yes}, \text{CS: II2}, \text{HP: Absent}, (T) \rangle$	1
4	T	142
5	$\langle \text{HC: High-grade}, \text{CP: Perforation}, \text{GHS: Good}, \text{BD: No}, \text{CS: \{III, IV\}}, \text{HP: Present}, (T) \rangle$	4
6	$\neg T$	142

All the T explanations have CP: Perforation. The notation is the same as in Table 7. All the $\neg T$ explanations have been removed for the sake of clarity and because this table focuses on T .

indicate a new attribute importance ranking associated with the specific treatment.

For our chosen T , the resultant base is $B_T = [\text{HC}, \text{CP}, \text{GHS}, \text{BD}, \text{CS}, \text{HP}]$. Note that the importance of the last three attributes BD, CS, HP has changed. The $KBM2L$, shown in Table 8, has only 3 T -items (lengths 1, 1 and 4, respectively). The remaining cases correspond to $\neg T$. The resulting explanations are shown in bold face. All of them include CP=Perforation. A possible clinical interpretation of this result is as follows. As the patient presents with a perforated stomach, it is necessary to carry out curative surgery. This is combined with chemotherapy and radiotherapy, which is effective against low-grade lymphoma in the early stages, and also against high-grade lymphoma at a later stage. Antibiotics are not chosen, as the patient does not invariably have *H. pylori*.

Next, consider another treatment $T \equiv (\text{HT}=\text{Yes}, \text{S}=\text{Curative}, \text{CTRTS}=\text{Radio})$. It belongs to 6 items and the base is B_{final} , as shown in Table 6. Here, after relabeling the alternatives as T and $\neg T$, the resultant base does not change ($B_T = B_{\text{final}}$), and the number of T -items is still 6, see Table 9. In all items concerning treatment T , the selection of antibiotic treatment (HT=Yes) is based on the presence of the *H. pylori*. At the early stages low-grade gastric NHL is not particularly sensitive to chemotherapy, though radiotherapy is often effective in such patients. All patients have either a perforated stomach, gastric hemorrhage or gastric obstruction. Hence, curative surgery seems to be an appropriate choice.

5. Conclusions

It is difficult to reach consensus on treatment decisions in a field such as oncology, even if decisions are

Table 9

The optimal *KBM2L* list for treatment $T \equiv (\text{HT}=\text{Yes}, \text{S}=\text{Curative}, \text{CTRSTS}=\text{Radio})$ obtained with $B_T = [\text{HC}, \text{CP}, \text{GHS}, \text{HP}, \text{CS}, \text{BD}]$

# Item	Description	Size
0	$\neg T$	142
1	$\langle \text{HC: Low-grade, CP: Hemorrhage, GHS: Average, HP: Present, CS: \{I, II1\}, BD: \text{---}, (T) \rangle$	4
2	$\neg T$	142
3	$\langle \text{HC: Low-grade, CP: Perforation, GHS: Poor, HP: Present, CS: II2, BD: \text{---}, (T) \rangle$	2
4	$\neg T$	142
5	$\langle \text{HC: Low-grade, CP: Perforation, GHS: Average, HP: Present, CS: III1, BD: \text{---}, (T) \rangle$	2
6	$\neg T$	142
7	$\langle \text{HC: Low-grade, CP: Perforation, GHS: Average, HP: Present, CS: II2, BD: No, (T) \rangle$	1
8	$\neg T$	142
9	$\langle \text{HC: Low-grade, CP: Perforation, GHS: Good, HP: Present, CS: \{I, II1, II2\}, BD: \text{---}, (T) \rangle$	6
10	$\neg T$	142
11	$\langle \text{HC: Low-grade, CP: Obstruction, GHS: Average, HP: Present, CS: \{I, II1\}, BD: \text{---}, (T) \rangle$	3
12	$\neg T$	142

All the T explanations have HC: Low-grade and HP: Present. The notation is the same as in Table 8. All the $\neg T$ explanations have been removed for the sake of clarity and because this table focuses on T .

based on available evidence from medical science literature. Designing an influence diagram may help, as it forces one to specify the underlying reasons, including the uncertainties involved in the decision-making, to render the treatment decisions. Even if sound scientific evidence is missing, doctors still have to make decisions, and it is better to base the decisions on explicit knowledge rather than implicit. Such explicit decision-making knowledge may then be considered by clinical users in the context of their daily practice. The construction of an influence diagram yields such an explicit representation of the clinical knowledge involved in medical decision-making.

However, clinicians may find the resulting influence diagram hard to understand, as medical doctors are accustomed to a type of reasoning where decisions concerning clinical situations are based on clinical rules. Hence, an attractive alternative is to extract such rules from an already designed influence diagram and make medical doctors to visualize them. This may not only be useful for explanatory purposes, but also for validating an influence diagram. In this paper we have studied the potential of the *KBM2L* lists technique as an aid for explaining and understanding an expert-designed influence diagram concerning gastric NHL.

The regularity patterns regarding optimal treatments, which can be discovered in decision tables

depend not only on the problem concerned, but also on their internal organization. A good organization of such tables reduces the memory required for storage, but, more importantly in a medical setting, may also be useful for finding out which key attributes underlie the treatment recommendations. This is what we have investigated in the context of the gastric NHL influence diagram.

During the refinement of an influence diagram, medical experts involved in the construction process may study whether the generated explanations for the optimal treatment recommendations agree with their own knowledge. Based on these insights, parts of a diagram may then be improved.

An evaluation study of the utility of *KBM2L* lists would involve comparing the process of the design of an influence diagram, both aided and unaided, by exploiting *KBM2L* lists. This clearly would be a major undertaking, as it would necessarily be carried out by independent groups of developers with same levels of expertise in the domain to be modeled. Some recent papers provide criteria to evaluate the utility of the explanations like the length of the explanation and the confidence level [16]. As a preliminary evaluation of the technique, the generated optimal *KBM2L* lists for the influence diagram regarding gastric non-Hodgkin lymphoma were clinically interpreted by the third author, who has a background as a medical doctor. A more thorough evaluation study, where we will investigate the impact of *KBM2L* lists when used in the middle of the development of an influence diagram, will be undertaken in the near future. This will allow us to identify the nature and amount of changes made to the influence diagram, based on feedback obtained from *KBM2L* lists to the developers of the influence diagram.

Our research now targets implementing a method for performing sensitivity analysis within our framework.

Acknowledgments

Research supported by projects from CICYT (TSI2004-06801-C04-04, MTM2004-21099-E) and CAM (S-0505/TIC/0230). We are very grateful for the referees' useful suggestions.

References

- [1] C. Bielza, M. Gómez, S. Ríos-Insua, J.A. Fernández del Pozo, Structural, elicitation and computational issues faced when solving complex decision making problems with influence diagrams, *Computers and Operations Research* 27 (7–8) (2000) 725–740.
- [2] G.F. Cooper, A method for using belief networks as influence diagrams, in: R.D. Shachter, T.S. Levitt, L.N. Kanal, J.F. Lemmer

- (Eds.), Proceedings of the 4th Workshop on Uncertainty in Artificial Intelligence, Morgan Kaufmann, 1988, pp. 55–63.
- [3] R.O. Duda, P.E. Hart, D.G. Stork, *Pattern Classification*, 2nd edition, Wiley, New York, 2001.
- [4] S. Eidt, M. Stolte, R. Fishcer, *Helicobacter pylori* gastritis and primary gastric non-Hodgkin's lymphoma, *Journal of Clinical Pathology* 47 (1994) 436–439.
- [5] P.H. Farquhar, Utility assessment methods, *Management Science* 30 (1984) 1283–1300.
- [6] J.A. Fernández del Pozo, C. Bielza, M. Gómez, Knowledge organisation in a neonatal jaundice decision support system, in: J. Crespo, V. Maojo, F. Martín (Eds.), *Medical Data Analysis*, Lecture Notes in Computer Science, vol. 2199, Springer, Berlin, 2001, pp. 88–94.
- [7] J.A. Fernández del Pozo, C. Bielza, M. Gómez, A list-based compact representation for large decision tables management, *European Journal of Operational Research* 160 (2005) 638–662.
- [8] D.S. Johnson, C.H. Papadimitriou, Computational complexity, in: E.L. Lawler, J.K. Lenstra, A.H.G. Rinnooy, D.B. Shmoys (Eds.), *The Traveling Salesman Problem*, John Wiley and Sons, 1985, pp. 37–85.
- [9] M.S. Gönül, D. Önkül, M. Lawrence, The effects of structural characteristics of explanations on use of a DSS, *Decision Support Systems* 42 (3) (2006) 1481–1493. doi:org/10.1016/j.dss.2005.12.003.
- [10] D.E. Knuth, *The art of computer programming*, Fundamental Algorithms, vol. 1, Addison-Wesley, Reading, 1968.
- [11] R. Kohavi, Bottom-up induction of oblivious read-once decision graphs, in: F. Bergadano, L. De Raedt (Eds.), *Machine Learning: ECML-94*, Lecture Notes in Computer Science, vol. 784, Springer, Berlin, 1994, pp. 154–169.
- [12] C. Lacave, F.J. Díez, A review of explanation methods for heuristic expert systems, *Knowledge Engineering Review* 19 (2004) 133–146.
- [13] S. Lauritzen, D. Nilsson, Representing and solving decision problems with limited information, *Management Science* 47 (9) (2001) 1235–1251.
- [14] P. Lucas, H. Boot, B. Taal, Computer-based decision-support in the management of primary gastric non-Hodgkin lymphoma, *Methods of Information in Medicine* 37 (1998) 206–219.
- [15] P. Lucas, H. Boot, B. Taal, A decision-theoretic network approach to treatment management and prognosis, *Knowledge-based Systems* 11 (1998) 321–330.
- [16] B. Padmanabhan, A. Tuzhilin, Knowledge refinement based on the discovery of unexpected patterns in data mining, *Decision Support Systems* 33 (3) (2002) 309–321.
- [17] Z. Pawlak, Rough set approach to knowledge-based decision support, *European Journal of Operational Research* 99 (1997) 48–57.
- [18] R.D. Shachter, Evaluating influence diagrams, *Operations Research* 34 (6) (1986) 871–882.
- [19] M. Shaw, C. Subramaniam, G.W. Tan, M.E. Welge, Knowledge management and data mining for marketing, *Decision Support Systems* 31 (2001) 127–137.
- [20] M. Vomlelová, F.V. Jensen, An extension of lazy evaluation for influence diagrams avoiding redundant variables in the potentials, *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 12 (2004) 1–17.



Concha Bielza received her MS degree in Mathematics in 1989 from Complutense University, Madrid, and PhD in Computer Science in 1996 from Technical University of Madrid, Madrid. She is currently an Associate Professor of Statistics and Operation Research in the School of Computer Science at Technical University of Madrid.

Her research interests are primarily in the areas of probabilistic graphical models, decision analysis, metaheuristics for optimization, data mining, classification models and real applications. Her research has appeared in journals like *Management Science*, *Computers and Operations Research*, *Statistics and Computing*, *Naval Research Logistics*, *Journal of the Operational Research Society*, *Medical Decision Making* and as chapters of many books. Her teaching interests include Statistics, Simulation, Decision-Support Systems and Machine Learning.



Juan A. Fernandez del Pozo is currently Associate Professor of Statistics and Operations Research at School of Computer Science. He gained degree in Computer Science from Technical University of Madrid His research interest is decision analysis and is at present involved in the development intelligent decision-support systems based on influence diagrams and bayesian networks that perform knowledge acquisition in huge decision tables, knowledge

discovery and data mining on models output for explanation synthesis and sensitivity analysis. His articles have appeared in various academic journals including: *Springer Lecture Notes in Computer Science*, *Journal of Operational Research*, *Computers and Operations Research*, Kluwer Academic Publishers.



Peter J.F. Lucas is an associate professor with the Institute of Computing and Information Sciences at the Radboud University Nijmegen, the Netherlands. He has been involved in research in Artificial Intelligence, in particular knowledge-based systems, since the beginning of the 1980s. He has contributed to this area by theoretical as well as applied research, the latter for the major part focusing on the field of medicine. At the moment, his research interests include topics such as applied

logic and theorem proving, knowledge representation, decision-support systems, model-based diagnosis, Bayesian networks, and statistical machine learning. He has extensively published in computing science and medical informatics journals and conferences (more than 100 journal and conference publications), written and edited several books, organized a number of workshops in the field, and edited 6 thematic issues of journals on topics mentioned above.