



HAL
open science

Explaining Multicriteria Decision Making with Formal Concept Analysis

Alexandre Bazin, Miguel Couceiro, Marie-Dominique Devignes, Amedeo
Napoli

► **To cite this version:**

Alexandre Bazin, Miguel Couceiro, Marie-Dominique Devignes, Amedeo Napoli. Explaining Multicriteria Decision Making with Formal Concept Analysis. *Concept Lattices and Applications 2020*, Jun 2020, Tallinn, Estonia. hal-02909383v2

HAL Id: hal-02909383

<https://hal.science/hal-02909383v2>

Submitted on 28 Mar 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Explaining Multicriteria Decision Making with Formal Concept Analysis

Alexandre Bazin, Miguel Couceiro, Marie-Dominique Devignes, Amedeo Napoli

LORIA, Université de Lorraine – CNRS – Inria, Nancy, France
`firstname.name@loria.fr`

Abstract. Multicriteria decision making aims at helping a decision maker choose the best solutions among alternatives compared against multiple conflicting criteria. The reasons why an alternative is considered among the best are not always clearly explained. In this paper, we propose an approach that uses formal concept analysis and background knowledge on the criteria to explain the presence of alternatives on the Pareto front of a multicriteria decision problem.

Keywords: Multicriteria decision making · Explanation · Background knowledge.

1 Introduction

Decision makers are regularly faced with situations in which they have to choose a solution among a set of possible alternatives. These alternatives most often have to be evaluated against multiple conflicting criteria. A single best solution cannot always be identified and, instead, decision makers need to identify solutions that represent the best compromise between the criteria. Multicriteria decision making [2, 3] is a field that offers a number of methods aimed at helping decision makers in proposing a set of “best” solutions.

In this paper, we are interested in further helping decision makers by adding, to a set of solutions proposed by a multicriteria decision method, an explanation of the reasons supporting the choice of these solutions. We propose an approach based on formal concept analysis explaining the output of such a multicriteria decision method in terms of background knowledge related to the criteria involved in the multicriteria decision problem. To illustrate the proposed approach, we suppose that the decision maker computes the Pareto front of their decision problem as a multicriteria decision method. However, the approach is not restricted to explaining Pareto fronts but works for any method for which the presence of an alternative in the solutions is monotonic w.r.t. the set of considered criteria. To the best of our knowledge, this approach is the first to use formal concept analysis to provide an explanation of the output of multicriteria decision methods.

The paper is structured as follows. In Section 2, we present the necessary definitions of formal concept analysis and multicriteria decision making. In Section

3, we develop the proposed approach for explaining the output of a multicriteria decision method. In Section 4, we present the result of the approach to the problem of identifying the features that are the most efficiently used by a Naive Bayes classifier constructed on the public Lymphography dataset.

2 Definitions

In this section, we recall the notions of formal concept analysis and multicriteria decision making used throughout this paper.

2.1 Formal Concept Analysis

Formal concept analysis [5] is a mathematical framework, based on lattice theory, that aims at analysing data and building a concept lattice from binary datasets. Formal concept analysis formalises binary datasets as formal contexts.

Definition 1. (FORMAL CONTEXT)

A formal context is a triple $(\mathcal{G}, \mathcal{M}, \mathcal{R})$ in which \mathcal{G} is a set of objects, \mathcal{M} is a set of attributes and $\mathcal{R} \subseteq \mathcal{G} \times \mathcal{M}$ is a binary relation between objects and attributes. We say that an object g is described by an attribute m when $(g, m) \in \mathcal{R}$.

From a formal context, we define two derivation operators \cdot' such that

$$\begin{aligned} \cdot' : 2^{\mathcal{G}} &\mapsto 2^{\mathcal{M}} \\ G' &= \{m \in \mathcal{M} \mid \forall g \in G, (g, m) \in \mathcal{R}\} \\ \cdot' : 2^{\mathcal{M}} &\mapsto 2^{\mathcal{G}} \\ M' &= \{g \in \mathcal{G} \mid \forall m \in M, (g, m) \in \mathcal{R}\} \end{aligned}$$

Definition 2. (FORMAL CONCEPT)

Let $\mathcal{K} = (\mathcal{G}, \mathcal{M}, \mathcal{R})$ be a formal context. A pair $(G, M) \in 2^{\mathcal{G}} \times 2^{\mathcal{M}}$ is called a formal concept of \mathcal{K} if and only if $M = G'$ and $G = M'$. In this case, G is called the extent and M the intent of the concept.

Formal concepts represent classes of objects that can be found in the data, where the extent is the set of objects belonging to the class and the intent is the set of attributes that describes the class. The set of formal concepts in a formal context ordered by the inclusion relation on subsets of either attributes or objects forms a complete lattice called the *concept lattice* of the context.

Definition 3. (CLOSURE AND INTERIOR OPERATORS)

Let (S, \leq) be an ordered set and $f : S \mapsto S$ such that

$$s_1 \leq s_2 \Rightarrow f(s_1) \leq f(s_2)$$

and

$$f(f(s_1)) = f(s_1).$$

The function f is called a closure operator if $x \leq f(x)$ and an interior operator if $f(x) \leq x$.

The two \cdot' operators form a Galois connection and the compositions \cdot'' are closure operators.

Definition 4. (LECTIC ORDER) *Let \leq be a total order on the elements of \mathcal{G} and $X, Y \subseteq \mathcal{G}$. We call lectic order the partial order \leq_l on $2^{\mathcal{G}}$ such that $X \leq_l Y$ if and only if the smallest element in $X \Delta Y = (X \cup Y) \setminus (X \cap Y)$, according to \leq , is in X . We then say that X is lectically smaller than Y .*

For example, let us consider a five element set $\mathcal{G} = \{a, b, c, d, e\}$ such that $e \leq d \leq c \leq b \leq a$. The set $\{a, c, e\}$ is lectically smaller than the set $\{a, b, c\}$ because $\{a, c, e\}$ contains e , the smallest element of $\{a, c, e\} \Delta \{a, b, c\} = \{b, e\}$.

Definition 5. (IMPLICATIONS AND LOGICAL CLOSURE) *Let S be a set. An implication $[6, 4]$ between subsets of S is a rule of the form $A \rightarrow B$ where $A, B \subseteq S$. The logical closure $\mathcal{I}(X)$ of a set $X \subseteq S$ by an implication set \mathcal{I} is the smallest superset Y of X such that $(A \rightarrow B \in \mathcal{I} \text{ and } A \subseteq Y)$ implies $B \subseteq Y$.*

For example, let $\mathcal{I} = \{\{a\} \rightarrow \{b\}, \{b, c\} \rightarrow \{d\}\}$. The logical closure of $\{a, c\}$ by \mathcal{I} is $\mathcal{I}(\{a, c\}) = \{a, b, c, d\}$. The logical closure by an implication set is a closure operator.

2.2 Multicriteria Decision Making

Multicriteria decision making [2, 3] is a field that aims at helping a decision maker choose the “best” solutions among a set \mathcal{A} of alternatives evaluated against multiple conflicting criteria.

Definition 6. (CRITERION AND MULTICRITERIA DECISION PROBLEM) *A criterion c_i on the alternative set \mathcal{A} is a quasi-order (reflexive and transitive relation) \succsim_{c_i} on \mathcal{A} . We say that an alternative a_1 dominates (or is preferred to) an alternative a_2 for the criterion c_i when $a_1 \succsim_{c_i} a_2$. The pair $(\mathcal{C}, \mathcal{A})$, where \mathcal{C} is a set of criteria, is called a multicriteria decision problem.*

Criteria are rankings of alternatives. For example, cars can be ranked according to their price or their maximum speed and both rankings are not necessarily the same. Someone who wishes to buy a new car has to consider all the available cars (the alternatives) and compare them against their price and speed (the criteria).

A *decision method* is a function that, given a multicriteria decision problem, returns a set of alternatives considered to be “the best”. Solving the multicriteria decision problem $(\mathcal{C}, \mathcal{A})$ consists in obtaining a set of best alternatives through the application of a decision method. In the above example, a function that returns both the fastest and cheapest cars is a decision method.

Definition 7. (PARETO DOMINANCE) *Let $\mathcal{C} = \{c_1, \dots, c_n\}$ be a set of criteria and a_1 and a_2 be two alternatives. We say that a_1 Pareto-dominates a_2 , denoted by $a_1 \succ a_2$, when $a_1 \succsim_{c_i} a_2$ for all $i \in \{1, \dots, n\}$.*

Definition 8. (PARETO FRONT) *Let \mathcal{A} be a set of alternatives and \mathcal{C} be a set of criteria on \mathcal{A} . The Pareto front of the multicriteria decision problem $(\mathcal{C}, \mathcal{A})$, denoted by $Pareto(\mathcal{C}, \mathcal{A})$, is the set of alternatives that are not Pareto-dominated by any other alternative.*

The Pareto front is a set of alternatives that are “the best” according to the Pareto-dominance, which makes the computation of the Pareto front a decision method. For example, a car buyer can use the Pareto front to identify cars that are not both slower and more expensive than another. The Pareto front has the property that, for any two criteria sets C_1, C_2 such that $C_1 \subseteq C_2$, $a \in Pareto(C_1, \mathcal{A}) \Rightarrow a \in Pareto(C_2, \mathcal{A})$. Throughout this paper, we will use the Pareto front as our sole decision method. When the set of alternatives is clear from the context, we will call $Pareto(\mathcal{C}, \mathcal{A})$ the Pareto front of the criteria subset \mathcal{C} .

3 Explaining Multicriteria-based Decisions

3.1 The Problem

Let us suppose that we have a multicriteria decision problem, i.e. a set \mathcal{C} of criteria and a set \mathcal{A} of alternatives. We compute the Pareto front of this problem and we consider these alternatives to be the “best” alternatives w.r.t. the criterion set \mathcal{C} . We would like to know *why* these alternatives are the “best”, i.e. which criterion or set of criteria make these alternatives appear on the Pareto front. Is the alternative a on the Pareto front because it is ranked first by a particular criterion? Is a a good compromise between a subset of criteria? If so, do these criteria have a common characteristic that makes a a good choice?

As a running example, we consider the following situation. A decision tree was built using a dataset in which objects are described by five features $\{a_1, a_2, a_3, a_4, a_5\}$ and their membership to one of two classes called 0 and 1. The decision tree is used to assign classes to new, unlabelled objects in another dataset D that relies on the same features. The true classes of these new objects is known and the performance of the decision tree is related to its ability to assign the correct classes to objects. This performance can be quantified using measures [8]. Some of these measures are combinations of four values:

- True positive (TP), the number of objects belonging to the class 1 that have been assigned the class 1
- False positive (FP), the number of objects belonging to the class 0 that have been assigned the class 1
- False negative (FN), the number of objects belonging to the class 1 that have been assigned the class 0
- True negative (TN), the number of objects belonging to the class 0 that have been assigned the class 0

Some of the performance measures that rely on these four values are presented in Table 1. Let $m(D)$ denote the value of the measure m quantifying the performance of the decision tree when assigning classes to the objects of the dataset D . Let D_i^f denote the state of dataset D after a number i of random permutations of the values of the feature f . The impact of the feature f on the value of the measure m for the decision tree, denoted $Impact(f, m)$, is defined as the average variation of the value of m when the values that the feature takes in the dataset are randomly permuted, i.e., for a large enough number k of permutations,

$$Impact(f, m) \approx \sum_{i=1}^k \frac{m(D_i^f) - m(D)}{k}$$

A negative impact means that the decision tree performs worse when the feature is disturbed and thus that the decision tree relies on values of the features to achieve its performance as quantified by the measure. As impacts are numerical values, features can be ranked according to their impacts on the value of a measure.

	Real class 1	Real class 0		
Predicted class 1	True Positive (TP)	False Positive (FP)	Precision = $\frac{TP}{TP+FP}$	FDR = $\frac{FP}{TP+FP}$
Predicted class 0	False Negative (FN)	True Negative (TN)	FOR = $\frac{FN}{FN+TN}$	NPV = $\frac{TN}{FN+TN}$
	Sensitivity = $\frac{TP}{TP+FN}$	FPR = $\frac{FP}{FP+TN}$	FScore = $2 \frac{Precision \times Sensitivity}{Precision + Sensitivity}$	
	FNR = $\frac{FN}{TP+FN}$	Specificity = $\frac{TN}{FP+TN}$	Accuracy = $\frac{TP+TN}{TP+TN+FP+FN}$	
	Positive Likelihood Ratio = $\frac{Sensitivity}{FPR}$		MCC = $\frac{TP \times TN - FP \times FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}$	
	Negative Likelihood Ratio = $\frac{FNR}{Specificity}$			

Table 1. Some of the measures of a model’s performance based on the four values TP, FP, FN and TN.

Continuing our example, we would like to know which features are “the best” for the decision tree’s performance. As the rankings of the features induced by different performance measures are not necessarily identical, identifying “the best” features can be expressed as a multicriteria decision problem in which the alternatives are the features and the criteria are the rankings induced by the performance measures. In this example, we suppose that five measures are considered so the multicriteria decision problem involves five criteria $\mathcal{C} = \{Accuracy, Sensitivity, Specificity, FScore, Precision\}$ (later identified as c_1 to c_5) and five alternatives $\mathcal{A} = \{a_1, a_2, a_3, a_4, a_5\}$. We further suppose that the impacts of the features on the model’s score for the five measures translate into the following rankings:

$$\begin{aligned}
c_1 = \textit{Accuracy} : & a_1 \succ a_3 \succ a_4 \succ a_2 \succ a_5 \\
c_2 = \textit{Sensitivity} : & a_2 \succ a_3 \succ a_1 \succ a_4 \succ a_5 \\
c_3 = \textit{Specificity} : & a_1 \succ a_4 \succ a_3 \succ a_2 \succ a_5 \\
c_4 = \textit{FScore} : & a_2 \succ a_3 \succ a_1 \succ a_5 \succ a_4 \\
c_5 = \textit{Precision} : & a_2 \succ a_1 \succ a_4 \succ a_3 \succ a_5
\end{aligned}$$

Using the Pareto front as a multicriteria decision method, we obtain that three alternatives (features) are best, i.e. $Pareto(\mathcal{C}, \mathcal{A}) = \{a_1, a_2, a_3\}$. The alternatives a_4 and a_5 are Pareto-dominated by a_1 . It follows that the features a_1 , a_2 and a_3 are those that have the highest impact on the decision tree’s performance. In the next two sections, we develop our approach for explaining the presence of these alternatives on the Pareto front.

3.2 Identifying the Criteria Responsible for the Decisions

In order to explain why an alternative a is considered as the “best” in a multicriteria decision problem, we first have to identify the minimal criterion sets for which a is on the Pareto front. In other words, we are looking for an algorithm to compute the criterion sets C_1 such that $a \in Pareto(C_1, \mathcal{A})$ and, for all $C_2 \subset C_1$, $a \notin Pareto(C_2, \mathcal{A})$. As computing the Pareto front of a criterion set is computationally expensive, we have to navigate the powerset of criteria as efficiently as possible.

Let $\mathcal{F} = \{A \subseteq \mathcal{A} \mid \exists C \subseteq \mathcal{C}, A = Pareto(C, \mathcal{A})\}$ be the family of alternative sets A for which there exists a criterion set C such that $Pareto(C, \mathcal{A}) = A$. In our running example, $\mathcal{F} = \{\{a_1, a_2, a_3\}, \{a_1, a_2\}, \{a_1\}, \{a_2\}, \emptyset\}$. Multiple criterion sets can have the same Pareto front. Hence, the $Pareto(\cdot, \mathcal{A})$ function induces equivalence classes on the powerset of criteria: two criterion sets C_1 and C_2 are equivalent if and only if $Pareto(C_1, \mathcal{A}) = Pareto(C_2, \mathcal{A})$. For an alternative set $A \in \mathcal{F}$, we use $Crit(A)$ to denote the family of criterion sets C such that $Pareto(C, \mathcal{A}) = A$. In other words, $Crit(A)$ is the equivalence class of criterion sets for which A is the Pareto front. In our running example, $Crit(\{a_1, a_2\}) = \{\{c_1, c_5\}, \{c_3, c_5\}, \{c_1, c_3, c_5\}\}$. We want to compute the minimal elements of each equivalence class.

To do this, let us define the interior operator $i : 2^{\mathcal{C}} \mapsto 2^{\mathcal{C}}$ that, to a criterion set C , associates $i(C)$, the lexicographically smallest inclusion-minimal subset of C for which $Pareto(C, \mathcal{A}) = Pareto(i(C), \mathcal{A})$. The images of i are the minimal elements of each equivalence class. In our running example, the images of i are $\{c_1, c_2\}$, $\{c_1, c_4\}$, $\{c_2, c_3\}$, $\{c_3, c_4\}$ that have $\{a_1, a_2, a_3\}$ as Pareto front, $\{c_1, c_5\}$, $\{c_3, c_5\}$ that have $\{a_1, a_2\}$ as Pareto front, $\{c_1\}$, $\{c_3\}$ that have $\{a_1\}$ as Pareto front, $\{c_2\}$, $\{c_4\}$, $\{c_5\}$ that have $\{a_2\}$ as Pareto front and \emptyset that has an empty Pareto front.

Algorithm 1 is used to compute $i(C)$. It performs a depth-first search of the powerset of criteria, starting from C , to find the lexicographically smallest inclusion-minimal criterion set that has the same Pareto front as C . The algorithm attempts to remove the criteria in the set one by one in increasing order. A criterion

is removed if its removal does not change the Pareto front. As the criteria are removed in increasing order, it ensures the lexicographically smallest subset of C that has the same Pareto front is reached.

Algorithm 1: Algorithm for computing $i(C)$.

Input: Criterion set C , Alternative set \mathcal{A} and a total order on C
Output: $i(C)$

```

1 begin
2    $P \leftarrow \text{Pareto}(C, \mathcal{A});$ 
3    $R \leftarrow C;$ 
4   forall  $x \in C$  in increasing order do
5     if  $P == \text{Pareto}(C \setminus \{x\}, \mathcal{A})$  then
6        $R \leftarrow R \setminus \{x\};$ 
7 return  $R$ 
    
```

As interior and closure operators are dual, computing the images of i can be treated as the problem of computing closed sets. Algorithm 2 computes all the images of $i(\cdot)$ with a Next Closure-like approach [5, 1]. Note that, for a criterion set $C \subseteq \mathcal{C}$, we use \overline{C} to denote $\mathcal{C} \setminus C$.

Algorithm 2: Algorithm for computing all the images of $i(\cdot)$.

Input: Criterion set C , Alternative set \mathcal{A} and a total order on C
Output: The set of images of $i(\cdot)$

```

1 begin
2    $I \leftarrow \{\emptyset\};$ 
3    $C = i(\overline{\emptyset});$ 
4   while  $C \neq \mathcal{C}$  do
5      $I \leftarrow I \cup \{C\};$ 
6     forall  $x \notin C$  in decreasing order do
7        $D \leftarrow i(\overline{\{c \in C \mid c \leq x\} \cup \{x\}});$ 
8       if  $x == \min(C \Delta D)$  then
9          $C \leftarrow D;$ 
10 return  $I$ 
    
```

The minimal criterion sets for which an alternative a is on the Pareto front are the minimal elements of the equivalence classes $\text{Crit}(A)$ with A minimal such that $a \in A$. Therefore, not all the images of i are useful. We use

$$M(a) = \{i(C) \mid C \in \text{Crit}(A), a \in A \text{ s.t.}, \forall B \in \mathcal{F} \text{ s.t. } B \subset A, a \notin B\}$$

to denote the minimal sets of criteria for which the alternative a belongs to the Pareto front.

In our running example, the only element of \mathcal{F} containing the alternative a_3 is $\{a_1, a_2, a_3\}$. The images of i in $Crit(\{a_1, a_2, a_3\})$ are $\{c_1, c_2\}$, $\{c_1, c_4\}$, $\{c_2, c_3\}$, $\{c_3, c_4\}$. These are the minimal criteria sets for which a_3 appears on the Pareto front. Overall, we have

$$\begin{aligned} M(a_1) &= \{\{c_1\}, \{c_3\}\} \\ M(a_2) &= \{\{c_2\}, \{c_4\}, \{c_5\}\} \\ M(a_3) &= \{\{c_1, c_2\}, \{c_1, c_4\}, \{c_2, c_3\}, \{c_3, c_4\}\} \end{aligned}$$

Note that alternatives a_1 and a_2 are on the Pareto front because of single criteria while a_3 requires multiple criteria to be on the Pareto front.

3.3 Explaining the Decisions

Once we have identified the minimal sets of criteria for which an alternative is on the Pareto front, we interpret them using background knowledge. This background knowledge takes the form of a set \mathcal{T} of terms and a set \mathcal{B} of implications between terms. In our running example, we arbitrarily choose to use as terms the names of the performance measures as well as TP , FP , TN , FN and $Measure$. The name of a measure implies the names of the values used in the measure's numerator. Hence, $\{Sensitivity\} \rightarrow \{TP\}$ because the TP value is used in $Sensitivity$'s numerator. Similarly, $\{FScore\} \rightarrow \{TP, FN, FP\}$ because both $Precision = \frac{TP}{TP+FN}$ and $Sensitivity = \frac{TP}{TP+FP}$ are used in $FScore$'s numerator. Additionally, all the terms imply $Measure$ because both measures and values used in measures are measures. This background knowledge induces the partial order on the set of terms depicted in Fig. 1.




Fig. 1. Example of background knowledge as a partially ordered set of terms.

Considering this partial ordering as an interpretation domain, we assign to a criterion c_i an interpretation by means of a *criterion interpretation function*

$\mathcal{I}_c(\cdot) : \mathcal{C} \mapsto 2^{\mathcal{T}}$ that maps c_i to a subset of \mathcal{T} closed under the logical closure by the implications of \mathcal{B} . In our running example, $\mathcal{I}_c(c_i)$ is the logical closure of the name of the measure associated with the criterion c_i . Hence, since $c_2 = \text{Sensitivity}$, $\mathcal{I}_c(c_2) = \mathcal{B}(\{\text{Sensitivity}\}) = \{\text{Sensitivity}, TP, \text{Measure}\}$. Overall, in our running example:

$$\begin{aligned} \mathcal{I}_c(c_1) &= \mathcal{B}(\{\text{Accuracy}\}) = \{\text{Accuracy}, TP, TN, \text{Measure}\} \\ \mathcal{I}_c(c_2) &= \mathcal{B}(\{\text{Sensitivity}\}) = \{\text{Sensitivity}, TP, \text{Measure}\} \\ \mathcal{I}_c(c_3) &= \mathcal{B}(\{\text{Specificity}\}) = \{\text{Specificity}, TN, \text{Measure}\} \\ \mathcal{I}_c(c_4) &= \mathcal{B}(\{\text{FScore}\}) = \{\text{FScore}, FP, TP, FN, \text{Measure}\} \\ \mathcal{I}_c(c_5) &= \mathcal{B}(\{\text{Precision}\}) = \{\text{Precision}, TP, \text{Measure}\} \end{aligned}$$

We then use the interpretations of criteria to explain the reason why an alternative is on the Pareto front. In our running example, a_1 appears on the Pareto front as soon as either c_1 or c_3 is taken into account. As c_1 represents the impact of features on the accuracy score of the decision tree, we can say that a_1 is on the Pareto front because it is particularly good for accuracy. Similarly, as c_3 represents the impact of features on the specificity score of the decision tree, a_1 is also good for specificity. The alternative a_3 only starts appearing on the Pareto front when two criteria are taken into account simultaneously, such as c_1 and c_2 . In this situation, we can only say that a_3 is good for the commonalities of c_1 and c_2 . As c_2 represents the impact of features on the sensitivity score of the model, according to our background knowledge, we can say that a_3 is good for measures that use the TP value. Formally, we define the interpretation of an alternative by means of the *alternative interpretation function* $\mathcal{I}_a(\cdot) : \mathcal{A} \mapsto 2^{\mathcal{T}}$ that maps an alternative a_i to $\mathcal{I}_a(a_i) = \bigcup_{C \in M(a_i)} \bigcap_{c_j \in C} \mathcal{I}_c(c_j)$, i.e. a subset of \mathcal{T} closed under the logical closure by the implications of \mathcal{B} .

Finally, we use formal concept analysis to classify the alternatives according to their interpretations and present the explanation of the presence of alternatives on the Pareto front in the form of a concept lattice. The interpretation of an alternative provides its description, which gives rise to the formal context $(\text{Pareto}(\mathcal{C}, \mathcal{A}), \mathcal{T}, \mathcal{R}_a)$ in which $(a_i, t) \in \mathcal{R}_a$ if and only if $t \in \mathcal{I}_a(a_i)$. In our running example, this formal context is the one depicted in Table 2. In the concepts of this context, the extents are sets of alternatives on the Pareto front of the multicriteria decision problem and the intents are interpretations of these alternatives that explain the reason why these alternatives are on the Pareto front. Figure 2 presents the concept lattice of our running example. For legibility reasons, only the most specific terms, according to the background knowledge, are displayed.

From this concept lattice, we learn that a_2 is among the best features for the decision tree's performance because it is the best for the Sensitivity, FScore and Precision scores while a_1 is the best for the Accuracy and Specificity scores. The feature a_3 is not the best for any particular measure's score but is a good compromise for measures that rely on the TP value.

	Accuracy	Sensitivity	Specificity	FScore	Precision	TP	FN	FP	TN	Measure
a_1	×		×			×			×	×
a_2		×		×	×	×	×	×		×
a_3						×				×

Table 2. Formal context $(\{a_1, a_2, a_3\}, \mathcal{T}, \mathcal{R}_a)$ of our running example.




Fig. 2. Concept lattice presenting sets of alternatives from our running example together with an explanation of their presence on the Pareto front.

4 Experimental Example

The Lymphography dataset is a public dataset available on the UCI machine learning repository¹. In its binarised version, it contains 148 objects described by 68 binary attributes identified by numbers. We used it to train a Naive Bayes classifier and to rank the features according to their impacts on six measures: Sensitivity, Specificity, Precision, Accuracy, Fscore and NPV. Identifying the best features for the Naive Bayes classifier’s performance is a multicriteria decision problem in the same manner as our running example. Our method was applied to this multicriteria decision problem and produced the concept lattice depicted in Fig. 3.

In this concept lattice, we see that five features are the best for the model. The feature 62 is the best for the NPV, Sensitivity and FScore measures. The feature 46 is the best for the Specificity and Precision measures. The feature 66 represents a good compromise for measures that use the TP or TN value. The feature 63 is a good compromise for measures that use the TN values. Finally, the feature 19 is a good compromise for measures in general. This explanation can, for example, be used in a feature selection [7] process to optimise the performance of classifiers.

5 Conclusion

The proposed approach provides an explanation of the output of multicriteria decision methods. The approach requires the presence of an alternative in the

¹ <http://archive.ics.uci.edu/ml/datasets/Lymphography>




Fig. 3. Concept lattice presenting an explanation of the best features from the binarised Lymphography dataset according to a naive Bayes classifier.

solutions returned by the decision method to be monotonic w.r.t. the set of considered criteria. Decision methods that satisfy this property include, among others, the Pareto front and all methods that return ideals of the set of alternatives partially ordered by the Pareto dominance relation.

The explanations that are created rely on background knowledge of the criteria. This allows the approach to be applied to any problem in which rankings of elements are produced by processes that are understood in some way.

Acknowledgement

This work was supported partly by the french PIA project « Lorraine Université d’Excellence », reference ANR-15-IDEX-04-LUE.

References

1. Borchmann, D.: A generalized next-closure algorithm - enumerating semilattice elements from a generating set. In: Proceedings of The Ninth International Conference on Concept Lattices and Their Applications, Fuengirola (Málaga), Spain, October 11-14, 2012. pp. 9–20 (2012)
2. Bouyssou, D., Dubois, D., Prade, H., Pirlot, M.: Decision Making Process: Concepts and Methods. John Wiley & Sons (2013)
3. Bouyssou, D., Marchant, T., Pirlot, M., Tsoukias, A., Vincke, P.: Evaluation and Decision Models with Multiple Criteria: Stepping Stones for the Analyst, vol. 86. Springer Science & Business Media (2006)

4. Distel, F., Sertkaya, B.: On the complexity of enumerating pseudo-intents. *Discrete Applied Mathematics* **159**(6), 450–466 (2011)
5. Ganter, B., Wille, R.: *Formal Concept Analysis: Mathematical Foundations*. Springer - Verlag (1999)
6. Ryssel, U., Distel, F., Borchmann, D.: Fast algorithms for implication bases and attribute exploration using proper premises. *Annals of Mathematics and Artificial Intelligence* **70**(1-2), 25–53 (2014)
7. Saeys, Y., Inza, I., Larrañaga, P.: A review of feature selection techniques in bioinformatics. *Bioinformatics* **23**(19), 2507–2517 (2007)
8. Sokolova, M., Lapalme, G.: A systematic analysis of performance measures for classification tasks. *Information Processing & Management* **45**(4), 427–437 (2009)