What is a decision problem? Preliminary statements

Alberto Colorni[1] and Alexis Tsoukiàs[2]

[1] Design Department, Politecnico di Milano
[2] LAMSADE-CNRS, Université Paris Dauphine

Abstract. This paper presents a general framework about what is a decision problem. The aim is to provide a theory under which the existing methods and algorithms can be characterised, designed, chosen or justified. The framework shows that 5 features are necessary and sufficient in order to completely describe the whole set of existing methods. It also explains why optimisation remains the general approach under which decision problems are algorithmically considered.

1 Introduction

The reader should be aware that this paper does not address the title question in a comprehensive way. The problem of what is a decision and what is a decision problem has been addressed in philosophy, psychology and the cognitive sciences, economy, political science etc.. We are not going to make a survey of this literature which is out of the scope of the paper. The reader interested in these aspects can have a look to a number of fundamental texts such as [10], [12], [16], [24], [26], [28], [30], [35].

Our proposition is instead pretty technical and formal. Operational Research and Decision Analysis are seen as part of a more general Decision Aiding Methodology (see [32]) aiming to help real decision makers to understand, formulate and model their problems and possibly reach a reasonable solution (if any). We are concerned by that type of activities occurring in a decision aiding situation where a "client" (very broadly defined) asks for some advice or help to an "analyst", such an advice being expected to come under form of a formal model allowing some form of rationality. We call such activities a "decision aiding process" (see [31]). At a certain point of that process the analyst will have to formulate a "decision problem" requiring some computing to be performed by some algorithms providing a result which is expected to be used in order to present a recommendation relevant to the decision maker's "decision problem".

Our focus is exactly here: what is a decision problem for the analyst? The proposal of the paper is to suggest a general framework under within which it is possible to identify all possible models, algorithms, procedures which routinely analysts use in their job as well as to allow to invent ones (if possible). The paper introduces two hypotheses:
- It is possible to establish a common framework under which any formal decision problem can be formulated, enabling to construct wide classes of methods characterised by common features.
- From an algorithmic point of view any decision problem can be reduced to an optimisation problem.

In the following section we introduce notation. Then in section 3 we show what the primitives of a decision problem are. Then in section 4 we describe the five characteristic features under which methods can be described. Section 5 introduces some methodology principles, while section 6 discusses the two running examples of the paper. Further research challenges are introduced within the conclusion.

2 Concepts and Notation

In the following A will always represent the set of "alternatives" considered either within a model or by a method. Although in practice such a set is never readily available, but constructed, for the purpose of this paper we are going to consider it as "given".

Along the paper we are going to use extensively preference relations. The basic relation we will adopt will be $\succeq$ (possibly indexed $\succeq_i$) which will read as "at least as good as" ($\succ$ will represent the asymmetric part of $\succeq$, while $\sim$ will represent the symmetric part). We will only make the hypothesis that this is a reflexive binary relation. The interested reader can see more about preference structures in [21] or [27] from which we adopt definitions and notation. We are now able to make our first claim.

Claim 1 *A decision problem for the analyst consists in finding an appropriate partitioning of the set A, relevant for the decision maker's concerns.*

The presentation of any algorithm or method discussed in this paper will be based on separating the "primitives" (what is the strictly necessary information required to be provided by the decision maker in order to allow some reasonable advice) and the "output" (what is the information the algorithm or method provides to the user). The reader should note that we distinguish between primitives and "input" to the algorithm. The reason is that the input to a precise existing algorithm or method is in reality constructed out of the primitives. Let's summarise. Our hypothesis is that the modelling process, that is the dialogue between the client and the analyst, follows (roughly) a sequence starting with the client providing *ground information*, which through *learning protocols* is transformed in *primitives* and these through *modelling tools* are transformed to the input to some method. For the description of these concepts see figure 1.

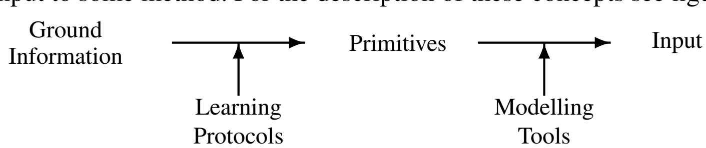

Fig. 1. The modelling process

Ground Information contains the problem description and for the purposes of this paper we will focus to what we call "preference statements": pieces of client's statements (in his own language), expressing values, opinions and likelihoods. In other terms it is how the client see his/her problem. *Learning Protocols* are procedures allowing to identify preference statements within the client's discourse and to translate them in ordering relations. In order to do so we need to establish the sets on which such relations apply. As it will become clear in section 4 such protocols are aimed at establishing the set A, the problem statement and the preference relations upon A. *Primitives* are the ordering relations "learned" using the protocols and we will discuss them extensively in the next section. *Modelling Tools* are the usual analytic tools an analyst uses in order to transform primitives in decision aiding models. Examples include the procedures allowing to construct a value function, a set of constraints, a probability distribution etc.. The *Input* is the information modelled in such a way that a decision aiding method can be applied. For instance in a linear programming method the input are the decision variables, the constraints and the objective function. In the following we present two running examples explaining some of the concepts introduced in each section.

Example 1.1 *Ground information*. The client is a horse races gambler. He is considering the next bet to make. In order to assess the "value" of each possible bet the client considers three different information: the quality of the horse, the quality of the jockey who runs it and the weather conditions. The client wants to rank all possible bets.

Example 2.1 *Ground information*. A hospital is considering the recrutement of nurses for three of their departments: General Medicine (GM), Oncology (ON), Children (CH). The hirings are managed by two: the general manager and the surgeon general. Candidates fill an application form and go through an interview. Practically the result is a report where the two managers consider three information: the age, the specialisation (if any) and the motivations of the candidate.

In both cases the learning protocols are procedures through which the analyst will try to gather the preferences of the client(s). Which horses (s)he prefers? With which jockey? Under which weather conditions? What is a good nurse for a given department? How specialisations compare with respect to the requirements of each department? How age influences the fitting of a candidate to a given department?

3 Primitives and Problems

What type of information can generally affect a decision? Since the origins (see for instance [7], [23]) most of the decision analysis literature will classify such information in three categories: values, opinions and likelihoods.

1. Values (related to attributes). Values should represent "what matters for the decision maker" (for a nice discussion see [13]). Under a more formal perspective we consider that the set A can be described against a set of attributes D, each attribute being equipped with a scale from a set of scales E. Following measurement theory (see [25]) such scales can be nominal, ordinal, ratio or interval ones. However, this is just descriptive information about A (x is 10cm long, y is yellow etc.). In order to be able to talk about values affecting decisions we need further information coming under form of preferential statements ("I prefer long tables to short ones", " I do not like yellow shoes"), possibly of more complex content ("I prefer a train travel to Paris to a flight at Amsterdam", "my preference of apples against oranges is stronger than my preference of peaches against apricots").We distinguish two types of sentences:
- *comparative ones*, where an elements of A are compared among them (under one ore more attributes) in order to express a preference;

- *absolute ones*, where an element of A is directly assessed against some "value structure" (under one ore more attributes).

2. Opinions (related to stakeholders). Decisions can be affected by the judgements and opinions of many stakeholders. In this case preference statements are going to be associated to "opinions". It is reasonable however, to distinguish once again among:
- *comparative opinions* (stakeholder i prefers x to y), where preferences are expressed among elements of the set A;
- *absolute opinions* (stakeholder i considers x as "worthy"), where preferences are expressed under form of value assessments.

3. Likelihoods (related to scenarios). When we express preferences it is likely that these depend from uncertain future conditions. Although the intuitive temptation is to use estimates (it is likely to rain) or quantifications of uncertainty (the probability of raining is p), if we focus on decision situations the primitives we need to consider will once again be preference statements of the type "under scenario j, I prefer x to y" or of the type "under scenario j, x is unworthy".

The reader will note that we do not include among the primitives the concept of relative importance of the dimensions under which preferences are expressed. The reason is simple: relative importance is a derivable information. Consider the case where $x \succ_I y$ and $y \succ_J x$ for some $I, J \subset H$ (x, y being elements of A and H being the set of criteria). If we add the information that "globally" $x \succ y$ (which is once again a primitive) we can derive that $I \gg J$ ($\gg$ representing an ordering relation upon the power set of H: $\gg \subset 2^H \times 2^H$). This will be true if I and J are sets of values, but also if they are opinions or likelihoods. Of course a decision maker may wish to make direct statements comparing two dimensions (I is more likely to occur than J), but there are two reasons for which it is better to avoid it. The first has to do with the fact that there is no general model for "relative importance", this depending on how primitive preferences are considered at the global level (see [17]). The second is that these are second order comparisons allowing for more cognitive biases and potential inconsistencies. Indeed more often than less decision makers are ready to change their statements about relative importance as soon as they realise the impact they may have on first order preference statements (that is comparing elements of A).

From the above discussion and in order to model "absolute statements" we need to introduce, besides the set A (of potential decisions), a set B being a collection of "norms" or "standards" or "thresholds" representing an external (with respect to A) value structure. In other terms if we want to claim that x is "nice" (under a certain point of view) we need to establish somewhere (not in A) what "nice" means and compare x to that norm. Under such a perspective:

Definition 1.
- *Comparative preference statements come under form of* $x \succeq_i y \;\; x, y \in A$*, i being any among attributes, stakeholders or scenarios (thus* $\forall i \succeq_i \subseteq A \times A$*).*
- *Absolute preference statements come under form of* $x \succeq_i b \;\; x \in A \; b \in B$*, i being any among attributes, stakeholders or scenarios (thus* $\forall i \succeq_i \subseteq A \times B \cup B \times A$*).*

From the above presentation we can establish our second general claim.

Claim 2 *Decisions are based on primitives which always come under form of comparative or absolute preference statements.*

Remark 1. We use the term preference statement in a very broad way. However, the formalism adopted should not conduct to confusion. Preference statements can be modelled either under form of asymmetric ordering relations (x is strictly before y) or under form of symmetric ordering relations such as similarities or nearness relations (x is similar or near to y) and these can be learned directly from the client. The use of the $\succeq$ relation is a comfortable way to combine such relations in a unique definition.

We continue with our running examples.

Example 1.2 *Primitives*. The alternatives considered by the client are the "bets" (a combination of a horse with its jockey). This is a finite enumeration of the participants to the next race. The client provides different types of preference statements (examples):
- horse x is better than horse y;
- jockey i is better than jockey j;
- horse w run by jockey j is better than horse z run by jockey i;
- if it rains horse y is better than horse w.

Such sentences need to be interpreted. For instance should we understand the first sentence as "horse x being better than horse y independently from the jockeys they run them and the weather conditions?" Or should we understand it as "considering the same jockey and the same weather conditions then horse x is better than horse y?" The difference can be important. In the first case we consider that $\forall j \in J, t \in T \langle xjt \rangle \succeq \langle yjt \rangle$ where J is the set of jockeys and T is the set of weather conditions. This will imply for instance that $\langle x, \text{Paul}, \text{rain} \rangle \succeq \langle y, \text{John}, \text{dry} \rangle$. In the second case such a comparison is not allowed. We can only write formulas of the type $\langle x, \text{Paul}, \text{rain} \rangle \succeq \langle y, \text{Paul}, \text{rain} \rangle$ (ceteris paribus comparisons).

Example 2.2 *Primitives*. The set A is composed by those candidates who filled the application form and got the interview. However, this is an assignment problem where we also need to specify the classes where the elements of A are assigned. These are four: the three hospital departments and the rejected. The type of preference statements we need here are of the type: "being young is ideal for the Childrens' Department", "not having a specialisation is very bad for Oncology", "being motivated and specialised are the candidates we are looking for", these being more or less nuanced among the two managers (who may possibly disagree on some of them).

However, once again these need to be better understood. For instance we need to establish what "young" means (we need a threshold for that). Perhaps too young is not that ideal, that implying measuring an absolute distance from some ideal "young nurse age". We also need to understand if not having a specialisation is an eliminating handicap for the candidate or if it is a general negative assessment becoming more important for the special case of Oncology. On the other hand the reader will note that at this stage we do not need to compare candidates among them.

The reader will note that in the case of the horse races we are using comparative preference statements (a horse is better than another), while in the case of the nurses we are using absolute preference statements (a nurse feature fits, more or less, the requirements of a given department).

Using claim 2 and definition 1 we can summarise the possible primitive information used in a decision problem as in figure 2. As we move away the origin along any of the axes we start considering multiple values (opinions, likelihoods).

Definition 2. *We call optimisation problem any decision problem considering a unique dimension under which primitives are expressed (point $\mathcal{O}$ of figure 2).*

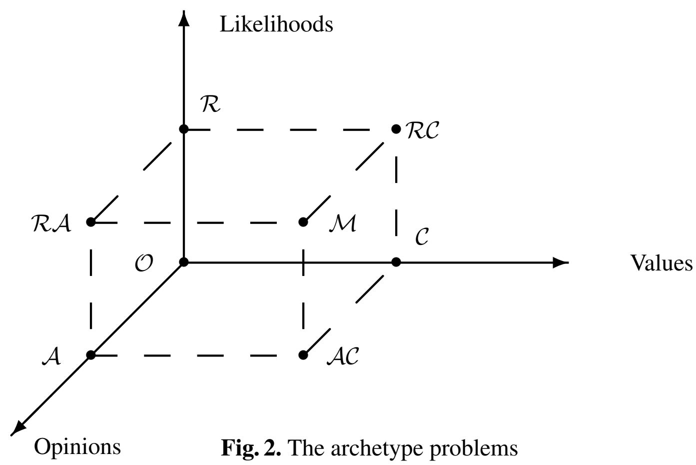

Fig. 2. The archetype problems

Figure 2 establishes eight archetype "problems" represented by the eight points: $\mathcal{O}, \mathcal{A}, \mathcal{C}, \mathcal{R}, \mathcal{AC}, \mathcal{RC}, \mathcal{RA}$ and $\mathcal{M}$. We call:

- $\mathcal{O}$: an optimisation problem (see more details in section 4);
- $\mathcal{A}$: an agreement problem, since different opinions need to be taken into account;
- $\mathcal{C}$: a compromise problem, since different values need to be considered;
- $\mathcal{R}$: a robustness problem, since a solution needs to be considered worthy under different likelihoods;
- $\mathcal{AC}$: an agreed compromise (a combination of $\mathcal{A}$ and $\mathcal{C}$);
- $\mathcal{RC}$: a robust compromise (a combination of $\mathcal{R}$ and $\mathcal{C}$);
- $\mathcal{RA}$: a robust agreement (a combination of $\mathcal{R}$ and $\mathcal{A}$);
- $\mathcal{M}$: a "mess", because the problem starts to become really messy ...

However, the above eight archetype problems do not stand alone. Behind a compromise problem other compromises may need to be considered in a hierarchy of criteria. Behind an agreement problem other agreement problems may have to be solved along a hierarchy of delegates, community representatives and other organisational structures. Behind a robustness problem many states of the nature may have to be considered in a hierarchy of likelihoods establishing complex scenarios. And, any combination of the above may in reality occur as complex as possible (see figure 3).

What can we observe in analysing these archetypal decision problems? Despite the different semantics behind values, opinions and likelihoods, the underlying formal structure is always the same. We have one common primitive: preference statements which we represent through preference relations. And we have one principal task: move along the hierarchy, from the leaves of the most elementary preference statements up to the root where "$x \succeq y$ all relevant information being considered" (x, y being either both in A or one in A and the other in B). This allows to introduce our third claim.

Claim 3 *A decision problem can be represented as a sequence of preference aggregations along an hierarchy of opinions, values and likelihoods, combined arbitrarily.*

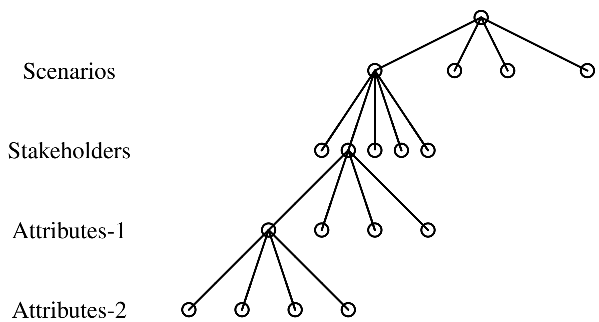

Fig. 3. Many decision problems hierarchically related

Example 3.1 Consider the case where a committee is assessing a number of development projects for an urban area. The outcomes of these projects depend on a number of uncertain issues due to the unstable economic situation of the whole region. At this stage we have many levels of the dimensions hierarchy: the different members of the committee and the different scenarios considered as "realistic" for the region, the set of attributes describing the projects. We consider as first level of our hierarchy the scenarios and as second level the committee members. We can expect that each committee members will assess the projects on a number of attributes (such as cost, sustainability, environmental impact etc.). It is reasonable to consider that some of such attributes decompose further in other attributes (such as direct costs, maintenance costs, financial costs etc.). This situation is captured by figure 3.

4 Main Features

In the following we present the 5 features which constitute the key parameters designing the whole set of conceivable formal decision problems. For the time being only one feature (the problem statement) will be discussed in an extensive way, the other four being essentially sketched.

4.1 The set of alternatives

A can be of different types:
- a countable enumeration of objects, $A = \{a_1 \cdots a_n\}$;

- a subset of all possible combinations of the attribute scales in the attributes space, $A \subseteq E_1 \times \cdots E_m$;
- a combinatorial structure resulting from the product of a set of discrete decision variables (possibly 0 or 1), $A \subseteq X_1 \times \cdots X_m, \ \ X_i \subseteq \mathbb{Z}$ or $X_i = \{0, 1\}$;
- a vector space resulting from the product of a set of real valued decision variables, $A \subseteq X_1 \times \cdots X_m, \ \ X_i \subseteq \mathbb{R}$.

4.2 The Problem Statements

Partitioning a set A consists in establishing a set of "equivalence classes" to which associate the elements of A. We can distinguish two different cases:

1. The first case concerns the fact that such classes can be ordered or not. Typical examples in the first case are classes of merit, the equivalence classes of a weak order etc.. Typical examples of the second case are problems of medical diagnosis, failure detection, pattern recognition etc..

2. The second case concerns the fact that such classes can be pre-defined with respect to some norm, standard, profile etc. or not. Typical examples of the first case include assigning elements of A to given ratings or patterns. Typical examples of the second case are clustering a population for some attribute or ranking it.

We summarise the above cases in table 1.

	Pre-defined wrt some external standard	NOT pre-defined
Ordered	Rating	Ranking
Not Ordered	Assignment	Clustering

Table 1. Basic Problem Statements

A special cases within the above problem statements is the one where the number of classes are just two, one being the complement of the other.
Let's discuss more in details the above problem statements.

1. Ranking. The primitive in this case will be a binary relation on A: $\succeq \subseteq A \times A$ to be read "at least as good as". The expected result is a partitioning of A in equivalence classes $[A]_1, \cdots [A]_n$ such that:
• $\exists \succcurlyeq \subset A \times A$
• $\succcurlyeq \ = \ \succ \cup \approx$
• $[A]$ is the set of equivalence classes constructed by $\approx$
• $\succ \subset [A] \times [A]$, $\succ$ being a strict partial order such that:
• $[A]_j \succ [A]_i \ \Leftrightarrow \ j > i$ and
• $\forall x \in [A]_j, y \in [A]_i : x \succ y$ and
• $\forall x, y \in [A]_j \ \ x \approx y$

Discussion. The reader will note that the ordering relation among the equivalence classes is not the primitive relation comparing the elements of A. Generally speaking $\succeq$ is not an ordering relation since preferences can be partial and or inconsistent. If we have to proceed with some operational procedure we need to transform the preference relation $\succeq$ to an ordering relation $\succcurlyeq$. We may impose that $\succeq \subset \succcurlyeq$, but this is not mandatory (for instance in case of inconsistent preference statements we may want to drop some primitive comparisons). As already mentioned $\succ$ is not necessary a complete relation. In case we impose completeness ($\succ$ becoming a total order) we get that $[A]_1 = \sup_A(\succ)$ (n being the number of equivalence classes in which A is partitioned).

What is a choice problem? We consider as choice the particular case where A is partitioned in two classes $[A]_1 \succ [A]_2$. We can generalise this observation claiming that any ranking problem partitions A in a set of classes, the first being the "optimal elements", then the second one being the "second optimal ones" and so on. The current literature considers as optimisation the special case where:
$\succeq = \succcurlyeq$ and $\succcurlyeq$ is a weak order on A such that $\exists f : A \mapsto \mathbb{R} : x \succcurlyeq y \;\Leftrightarrow\; f(x) \geq f(y)$.
Under such a hypothesis it is clear that $[A]_1 = \max_A f(x)$. However, we do not really need the function f in order to "optimise". Generalising the concept of optimisation we can always construct an algorithm such that $[A]_1 = \sup_A(\succcurlyeq)$. Extending further this reasoning the whole set of equivalence classes can be constructed as result of some optimisation: $[A]_{n+1} = \sup_{A \setminus [A]_n}(\succcurlyeq)$ etc..

2. Clustering The primitive in this case is a binary relations on A: $\approx \subseteq A \times A$ to be read "similar to". The expected result is a partitioning of A in $[A]_1, \cdots [A]_n$ such that: $\forall x, y \in [A]_j \;\; x \approx y$ and $\forall x \in [A]_j, \; y \in [A]_i : \neg(x \approx y)$.

Discussion. In case $\approx$ is an equivalence relations then the partitioning of A results in constructing the indiscernibility relation on A ([22]). However, this is not generally the case. Elements of A are more or less similar between them. (or differently similar, see [34]). Under such a perspective we consider that instead of a single similarity relation we have a set of nested similarity relations $\approx_l$ and $[A]_j = \sup_A(\approx_l)$. In other terms we try to maximise similarity within classes (clusters) and minimise similarity among classes (clusters). If $\approx_l$ are nested similarity relations with nice properties then we can establish metrics (see [11]:
- $s(x, y)$: how similar is x to y?
- $d(x, y)$: how distant is x from y?

Then, establishing the equivalence class of any element $y \in A$,
$[A]_y = \{x | \max_A F(s(x, y))\}$, F being a measure (a fitting function) of the overall similarity of the elements of $[A]$ with respect to y, we can construct the clusters $[A]_j$. Meyer and Olteanu generalised this idea (see [18]) for general preference structures.

Remark 2. The reader should note that both ranking and clustering problem statements boil down in solving some mathematical optimisation problem. This should not be surprising: in absence of any external information and being allowed only to compare elements of A among them, the only mathematical notion we have, in order to clearly separate classes between them, is the one of "optimality".

3. Rating The primitive here is a binary relation from the set A to the set B: $\succeq \subseteq A \times B \cup B \times A$ to be read "at least as good as", B being the set of external "norms" characterising the ordered classes $C_1 \rhd \cdots \rhd C_n$. The expected result is to assign each element of A in a C_j such that: $x \in C_j \;\Leftrightarrow\; x \succcurlyeq p_j, p_{j+1}, \cdots p_n$ and $p_1 \cdots p_{j-1} \succcurlyeq x$.

Discussion. As in the ranking problem statement we need to differentiate between the primitive preference relation $\succeq$ and the operational result represented by the ordering relation $\succcurlyeq$. By transitivity of $\succcurlyeq$ it is clear that if element x is in C_j and element

y is in C_{j+1}, $x \succ y$, while if both elements are assigned in the same class $x \sim y$ ($\succ \cup \sim = \succcurlyeq$). However, the reader should remember that classes are pre-established.

Suppose now that the relation $\succcurlyeq$ is a weak order such that we can establish a function $f : A \cup B \mapsto \mathbb{R}$ such that $x \succcurlyeq p_j \Leftrightarrow f(x) \geq f(p_j)$. The problem of assigning the elements of A to the ordered classes represented by the "norms" p_j turns to be a classical constraint satisfaction problem: $C_j = \{x : f(x) \geq f(p_j)$ and $f(p_{j-1}) \geq f(x)\}$. We can generalise this concept dropping the function and claim that the rating problem can be considered a generalised constraint satisfaction problem.

4. Assignment The primitive is a binary relation on A: $\approx \subseteq A \times B \cup B \times A$ to be read "similar to", B being the set of external "norms" characterising the classes $C_1 \cdots C_n$ (the difference with the rating problem statement being the fact that these classes are not ordered among them). The result is to assign each element of A in a C_j such that: $x \in C_j \Leftrightarrow x \approx p_j$, where p_j is the norm characterising class C_j.

Discussion. Assigning objects to unordered classes could be seen as a constraint satisfaction problem where constraints are expressed as equalities ($C_j = \{x : f(x) = f(p_j)\}$), where f is a function representing a metric of similarity.

Concluding: the problem statements we present here can be handled either as an optimisation problem or as a constraint satisfaction one. Considering that any constraint satisfaction problem can be transformed in an optimisation one we can state one of our principal claims, based on the hypothesis that our problem statements are exhaustive of all possible decision problems (partitionings).

Claim 4 *From an algorithmic point of view any decision problem is an optimisation problem.*

Remark 3. The reader should not make confusion with the notion of decision problem typical in the algorithmic complexity literature ([8]). On the other hand we want to emphasise that what we are talking here concerns how algorithmically primitives get transformed in ordering relations such that can be used for recommending something or being used for further aggregations.

4.3 Independence

As already mentioned primitives come under form of statements of the type "x is at least as good as y, under I", x, y being mono or multi-dimensional objects and I being a subset among values, likelihoods and opinions. However, despite its intuitive meaning, such a sentence can still be interpreted in different ways. We are going to distinguish two principal interpretations:

- "x is at least as good as y, under I", independently on what happens to $H \setminus I$ (H being the set of criteria);
- "x is at least as good as y, under I", provided a condition holds in some $J \subseteq H \setminus I$.

These two interpretations lead to completely different problem formulations and consequently to different methods and resolution algorithms. Preferential independence (the first interpretation) allows to envisage a linear (additive) model representing preferences. Conditional preferences lead to more complex preference structures implying non linear aggregation functions ([9], [14], [29]) or specialised algorithms ([3], [4]).

4.4 Differences of Preferences

Let's recall once again our primitives and let's consider the sentence "x is strictly better than y and these are both better than z (under I)". We know we can represent this sentence giving numerical values to x, y, z (for instance $x = 3, y = 2, z = 1$ and adopt the natural ordering of the numbers. However we could choose the numerical representation $x = 100, y = 10, z = 1$ and it would be the same. Preferences are orders and the numbers we use only carry ordinal information.

The point is that in many cases we could either have richer information (we know for instance that x is twice more heavy than y) or we would like to have richer information of the type "x is much more better than y". We need to reason in terms of "differences of preferences" and their representation. In other terms we need primitives of the type: "xy is not less than zw" where xy (sw) represents the difference of preference between x and y (z and w). It is interesting to note that primitives of this type can be used also in order to express ordinal preferences, while the opposite is not true. Under such a perspective we can claim that primitives should always be considered as sentences about differences of preferences, the ordinal case being a special one. The interested reader can see more in the literature about conjoint measurement (see [15]) and how this helped in reframing multidimensional preferences (see [1], [2], [17]).

4.5 Positive and Negative Reasons

Consider a preference statement of the type: "I do not like x", or "any candidate, but not x". Such statements can be considered as explicit "negative preferential statements" to be considered independently from the "positive ones" (which are the usual ones). The idea here is that there are cases where decision makers need to express negative judgements and values which are not complementary to the positive ones (such as a veto on a specific dimension). Such statements have been explicitly considered in the literature both in decision theory (see [6], [20], [33]) and in argumentation theory (see [19]). When such situations occur we need to develop specific procedures adding thus a further dimension of characterisation of the decision problem at hand.

5 Methodology

Let's summarise in order to outline how our framework can be used for methodological purposes. We have a set A, information describing the set A against a number of attributes and preference statements (these being values, opinions or likelihoods) comparing either the elements of A among them or the elements of A to elements of a set B (the set of norms). We aim at partitioning A appropriately.

The first problem we have is reducing the problem to an "optimisation problem": that is, obtaining one-dimensional preference statements. In other terms we are trying to aggregate preference statements expressed on several different dimensions to a single one. For the time we consider that transforming some attributes to "constraints" (thus bounding the space of feasible solutions) has already been considered in establishing the set A. How do the different parameters described in section 4 influence the design (or the choice) of an appropriate solution method?

Allowing to have explicit measures of differences of preferences allows to handle richer preferential information, such that we can consider to obtain at the aggregated level preference statements sufficiently rich to satisfy nice properties (for instance obtaining directly an ordering relation). In case we need to work with purely ordinal information we should expect the negative consequences of Arrow's impossibility theorem (see the discussion in [5]). In case preferential independence holds we are in the "easy case": given a set of primitives holding at the same level of the hierarchy of dimensions (see figure 3) these can be aggregated (possibly through a linear model) to the parent node. In case independence does not hold we have two options: either we need explicitly non linear models accounting for the observed dependencies or we need to reformulate the modelling dimensions (these options being not exclusive). This second case may result in aggregating more levels of the hierarchy in a single step.

The presence of explicit negative preference statements (not complementary to the positive ones) will result in duplicating the decision model creating an hierarchy of "negative reasons" (to be associated to the hierarchy of "positive reasons"). It will also require to establish how and when these two sources of information should merge.

The second problem we need to handle is to obtain, out of the one-dimension primitives computed in the previous step, an ordering relation allowing the partitioning of the set A (the recommendation to hand to the decision maker). It is clear that the type of problem statement adopted strongly influences how this step will be considered since it establishes both the type of primitives we need to construct and the type of algorithm to be used. Before concluding this discussion we note that the properties of the set A will also influence the design or the choice of the method for obvious algorithmic reasons.

At this stage is easy to show that the five features we introduced in this paper (properties of the set A (and B), problem statement, preferential independence, difference of preferences, explicit negative preference statements) are necessary in order to clearly establish, design and axiomatically characterise any model, algorithm and method aiming at handling a decision problem (on set A). Our claim, not demonstrated here, is that these features are also sufficient. We thus get:

Claim 5 *The properties of the set A, the type of problem statement, the holding or not of preferential independence, the explicit use of differences of preferences, the explicit use of negative preference statements, are the necessary and sufficient features for choosing, designing, justifying and axiomatically characterising any decision problem and the associated resolution methods and algorithms.*

Discussion. The reader should recall that the use of the "learning protocols" provides the analyst with some first basic information: the set A, the problem statement and the preferences about A. Without these information we cannot really make any tentative to formalise a decision problem of a client. However, there are many methods through which this information can be handled. In order to choose among them, to explain and justify them besides maintaining an axiomatic coherence) we need to know more about the preferences: whether preferential independence is met, whether differences of preferences hold or not and whether explicit negative statements are done. For instance the use of linear programming will make sense in a situation where the set A is a subset of a vector space, the problem statement is a ranking one (and more precisely a choice) and the preferences among the elements of A are such to allow a numerical representation which is also an interval scale of ”profits” (or costs), which means that differences of preferences hold, and this is additive, which means that preferential independence also holds, while there is no use of explicit negative statements.

6 Horse races and Nurses

Let's finalise the discussion of the two running examples.

Example 1.3 *Modelling*. Having established the correct primitives representing the client's preference statements we can check whether preferential independence is met and how differences of preferences are measured. If the preferences among horses are independent from the preferences among jockeys and among weather conditions then we can clearly look for a linear aggregation of such values (in case differences of preferences are meaningful) and even allow to compute a relative importance for each weather condition under form of probabilities. This will allow to use expected utility theory.

Discussion. The problem presented is a “robust compromise”. A compromise because we need to take into account two different type of values (the quality of the horse and the quality of the jockey) and a robustness problem because we need to take into account three different likelihoods: raining, humid weather, dry weather. The expected utility problem formulation will hold under very specific conditions. In case these do not hold we can look for alternative ones (using CP nets [3], [4] or fuzzy integrals [9]).

Example 2.3 *Modelling*. Each candidate is described against three attributes which are very different: age is a continuous numerical scale (probably discretised), specialisation is a nominal scale (which specialisation, if any, has the candidate), motivation is (probably) an ordinal scale reporting the judgement of the expert who did the interview. On the other hand the classes also need to be described against the same attributes either providing the “ideal” values for each class (for instance the “ideal nurse” for the Childrens' Department is 30 year old) or the “minimal” ones (not more than 35 years old). Then we need to establish how the candidates compare to such classes. For instance we need to value the distance between two ages or between two specialisations. The reader should remember that such models might be different among the two managers. Should we be able to measure these differences and should these be commensurable among them, then we can envisage to write a value model (possibly additive) assessing the “fitting” of each candidate for each category. Otherwise we could opt for some ordinal model using some majority principle (of the type: if two criteria agree that candidate x fits to class a then assign him/her there). However, the presence of negative assessments play an important role here, since they may exclude a candidate from a certain class independently from any other assessment (if the candidate does not have a specialisation in oncology cannot work in Oncology).

Discussion. The situation is a an Agreed Compromise since we need to find a compromise among the three fitness criteria and an agreement among the two managers. There are two paths leading to the final assignment. The first (more cooperative) consists in finding an agreement between the two managers for each criterion separately and then find a compromise using the agreed assessments. The second (more negotiation oriented) consists in finding the assignments for each manager and then try to find an agreement for the cases where these are different (possibly all of them). The reader will note that there is no reason for which these two paths lead to the same result.

7 Discussion and Conclusion

In this paper we introduced a general framework describing the whole set of methods and decision problems an analyst may have to design in order a problem provided by a client. We have shown that decision problems can be generally seen as optimisation problems after a complex hierarchy of values, opinions and likelihoods gets transformed to a single dimension through a sequence of "preference aggregations". In order to construct this general framework we make use of what we call primitives (the strictly necessary information to be provided in order to model meaningfully the decision problem). Then we described five main features which characterises any method aiming at modelling and solving the decision problem. Our general claim is that these five features (the type of the set A, the problem statement, the holding of preferential independence, how differences of preferences are considered, the presence of explicit negative reasons) are necessary and sufficient in order to choose, design, justify and characterise any decision support procedure. Further research includes on the one hand showing how formal argumentation theory can help in explaining and justifying methods design and their outcomes and on the other hand exploring how formal model reformulation techniques can help finding the most appropriate algorithm to adopt for resolution purposes.

References

1. D. Bouyssou, Th. Marchant, M. Pirlot, A. Tsoukiàs, and Ph. Vincke. *Evaluation and decision models with multiple criteria: Stepping stones for the analyst*. Springer Verlag, Boston, 1st edition, 2006.
2. D. Bouyssou and M. Pirlot. Conjoint measurement models for preference relations. In D. Bouyssou, D. Dubois, M. Pirlot, and H. Prade, editors, *Decision Making Process*, pages 617 – 672. J. Wiley, Chichester, 2009.
3. R. Brafman and Y. Dimopoulos. Extended semantics and optimization algorithms for CP-networks. *Computational Intelligence*, 20(2):219 – 245, 2004.
4. R. Brafman and C. Domshlak. Preference handling: An introductory tutorial. *AI Magazine*, 30(1):58–86, 2008.
5. D. Dubois, H. Fargier, P. Perny, and H. Prade. Qualitative decision theory: from Savage's axioms to non-monotonic reasoning. *Journal of the ACM*, 49:455–495, 2002.
6. D. Dubois and H. Prade. An introduction to bipolar representations of information and preference. *International Journal of Intelligent Systems*, 23:866–877, 2008.
7. S. French. *Decision theory - An introduction to the mathematics of rationality*. Ellis Horwood, Chichester, 1988.
8. M. Garey and D. Johnson. *Computers and Intractability*. Freeman and Co., New York, 1979.
9. M. Grabisch and C. Labreuche. Fuzzy measures and integrals in MCDA. In J. Figueira, S. Greco, and M. Ehrgott, editors, *Multiple Criteria Decision Analysis: State of the Art Surveys*, pages 563–608. Springer Verlag, Boston, Dordrecht, London, 2005.
10. J. Habermas. *Theorie des kommunikativen Handelns*. Suhrkamp, Frankfurt am Main, 1981. engl. version: The Theory of Communicative Action, trans. by Thomas McCarthy, Cambridge: Polity, 1984-87.

11. M.F. Janowitz. *Ordinal and Relational Clustering*. World Scientific, Singapore, 2010.
12. M. Kaplan. *Decision Theory as Philosophy*. The Cambridge University Press, Cambridge, 1996.
13. R.L. Keeney. *Value-Focused Thinking. A Path to Creative Decision Making*. Harvard University Press, Cambridge, 1992.
14. R.L. Keeney and H. Raiffa. *Decisions with multiple objectives: Preferences and value trade-offs*. J. Wiley, New York, 1976.
15. D.H. Krantz, R.D. Luce, P. Suppes, and A. Tversky. *Foundations of measurement*, volume 1: Additive and polynomial representations. Academic Press, New York, 1971.
16. R.D. Luce and H. Raiffa. *Games and Decisions*. J. Wiley, New York, 1957.
17. Th. Marchant. Towards a theory of MCDM: stepping away from social choice theory. *Mathematical Social Sciences*, 45:343–363, 2003.
18. P. Meyer and A.-L. Olteanu. Formalizing and solving the problem of clustering in mcda. *European Journal of Operational Research*, 227:494–502, 2013.
19. S. Modgil and H. Prakken. A general account of argumentation with preferences. *Artificial Intelligence*, 195:361–397, 2013.
20. M. Öztürk and A. Tsoukiàs. Bipolar preference modelling and aggregation in decision support. *International Journal of Intelligent Systems*, 23:970–984, 2008.
21. M. Öztürk, A. Tsoukiàs, and Ph Vincke. Preference modelling. In M. Ehrgott, S. Greco, and J. Figueira, editors, *State of the Art in Multiple Criteria Decision Analysis*, pages 27 – 72. Springer Verlag, Berlin, 2005.
22. Z. Pawlak. *Rough Sets - Theoretical Aspects of Reasoning about Data*. Kluwer Academic, Dordrecht, 1991.
23. H. Raiffa. *Decision Analysis – Introductory Lectures on Choices under Uncertainty*. Addison-Wesley, Reading, MA, 1968.
24. N. Rescher. *Introduction to Value Theory*. Prentice Hall, Englewood Cliffs, 1969.
25. F.S. Roberts. *Measurement theory, with applications to Decision Making, Utility and the Social Sciences*. Addison-Wesley, Boston, 1979.
26. J. Rosenhead. *Rational analysis of a problematic world*. J. Wiley, New York, 1989. 2nd revised edition in 2001.
27. M. Roubens and Ph. Vincke. *Preference Modeling*. LNEMS 250, Springer Verlag, Berlin, 1985.
28. H.A. Simon. *Administrative behaviour: a study of Decision Making Processes in Administrative Organizations*. Mac Millan, New York, 1947.
29. P. Suppes, D. H. Krantz, R.D. Luce, and A. Tversky. *Foundations of Measurement*, volume 2: Geometrical, threshold and probabilistic representations. Academic Press, New York, 1989.
30. S. Toulmin. *The Uses of Argument*. Cambridge University Press, Cambridge, 1958.
31. A. Tsoukiàs. On the concept of decision aiding process. *Annals of Operations Research*, 154:3 – 27, 2007.
32. A. Tsoukiàs. From decision theory to decision aiding methodology. *European Journal of Operational Research*, 187:138 – 161, 2008.
33. A. Tsoukiàs, P. Perny, and Ph. Vincke. From concordance/discordance to the modelling of positive and negative reasons in decision aiding. In D. Bouyssou, E. Jacquet-Lagrèze, P. Perny, R. Slowinski, D. Vanderpooten, and Ph. Vincke, editors, *Aiding Decisions with Multiple Criteria: Essays in Honour of Bernard Roy*, pages 147–174. Kluwer Academic, Dordrecht, 2002.
34. A. Tversky. Features of similarity. *Psychological Review*, 84:327–352, 1977.
35. P. Watzlawick, J.H. Weakland, and R. Fisch. *Change; principles of problem formation and problem resolution*. Norton, New York, 1974.