Available online at www.sciencedirect.com

SCIENCE DIRECT®

European Journal of Operational Research 160 (2005) 741–761

EUROPEAN JOURNAL OF OPERATIONAL RESEARCH

www.elsevier.com/locate/dsw

Propositional distances and compact preference representation

Céline Lafage *, Jérôme Lang

IRIT—Université Paul Sabatier, 31062 Toulouse Cedex, France

Received 28 January 2002; accepted 15 June 2003
Available online 17 December 2003

Abstract

Distances between possible worlds play an important role in logic-based knowledge representation (especially in belief change, reasoning about action, belief merging and similarity-based reasoning). We show here how they can be used for representing in a compact and intuitive way the preference profile of an agent, following the principle that given a goal G, then the closer a world w to a model of G, the better w. We give an integrated logical framework for preference representation which handles weighted goals and distances to goals in a uniform way. Then we argue that the widely used Hamming distance (which merely counts the number of propositional symbols assigned a different value by two worlds) is generally too rudimentary and too syntax-sensitive to be suitable in real applications; therefore, we propose a new family of distances, based on Choquet integrals, in which the Hamming distance has a position very similar to that of the arithmetic mean in the class of Choquet integrals.

Keywords: Artificial intelligence; Propositional logic; Utility functions; Preference relations; Compact representations

1. Introduction

The specification of a decision making or a planning process necessarily includes the preferences of the agent. At the object level, the preference structure can be either a utility function assigning each possible consequence to a numerical value, or a weak order relation (possibly allowing for incomparibility in addition to indifference).

Now, once it has been specified what a preference structure is, the question of *how it should be represented*, or in other terms, how it should be encoded so as to be "computationally friendly", arises. A straightforward possibility consists in writing it down explicitly, namely by listing all possible consequences together with their utility values or by listing all pairs of consequences in the preference relation. Clearly, this explicit mode of representation is possible in practice only when the number of possible consequences is reasonably low with respect to the available computational resources.

* Corresponding author.
E-mail addresses: celine.lafage@yahoo.com (C. Lafage), lang@irit.fr (J. Lang).

0377-2217/$ - see front matter © 2003 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2003.06.037

The latter assumption often cannot be made, in particular when the set of alternatives has a combinatorial structure, that is, when the decision to make consists in giving a value to each of a given set decision variables: in this case, the set of alternatives is the Cartesian product of the value domains of the variables, and clearly, its cardinality grows exponentially with the number n of variables.

Consider for example that the agent has to express his preference about a dinner to be composed of a first course dish, a main course dish, a dessert and a wine, with a choice of 6 items for each. This makes 6^4 alternatives. This would not be a problem if the four items to choose were independent from the other ones: in this case, it would be satisfactory to represent independently the preferences about each of the variables; the joint preference over the set of dinners would then be determined using an aggregation function as it commonly done in multicriteria decision analysis: in our example, the preference over the set 6^4 alternatives would come down to four independent preference structures over sets of 6 alternatives each. Things become more complicated if the agent wishes to express *dependencies* between items, such as "I would like to have risotto ai funghi as first course, except if the main course is a vegetable curry, in which case I would prefer smoked salmon as first course", "I prefer white wine if one of the courses is fish and none is meat, red wine if one of the courses is meat and none is fish, and in the remaining cases I would like equally red or white wine", etc.

As a second example, consider the following problem concerning a decision to be made by a recruiting committee: when not a single, but k individuals (out of n) can be recruited, the space of possible alternatives can no longer be identified with the set of candidates, but has a combinatorial structure (it is the set of all possible sets of k candidates). A member of the committee can express her preferences explicitly only if the dependencies between individuals are ignored, which means that the members cannot express *correlations* between candidates, as for instance: "My favourite candidate is A, my second best is B and my third best is C; but since A and B work on similar subjects whereas C works on a complementary subject, the joint recruiting of A and C, or even of B and C, suits me better than the joint recruiting of A and B."

For such problems, the size of the alternative space and the impossibility to decompose the preference structure into smaller preference structures bearing on each single variable makes it unreasonable to ask agents to give (under the form of a table or a list) an *explicit* utility function or preference relation on the set of all solutions. Therefore, allowing for the expression of utility functions or preference relations over such sets of alternatives needs the definition of a language enabling to express the preference structure in a space as small as possible. Such languages, called *compact representation languages*, should also be as *expressive* as possible, and as close as possible to the intuition (ideally, the specification of a preference structure in the language should be easily follow from the agent's expression of his preferences in natural language). Lastly, these languages should be equipped with decision procedures, as efficient as possible, so as to enable the automation of the search for an optimal collective decision.

Such preference representation languages have been developed within the "knowledge representation" subcommunity of Artificial Intelligence. They are often build up on propositional logic, but not always (see for instance utility networks [1,19] or valued constraint satisfaction [26]). Logical languages allow for a concise representation of the preference structure, while preserving a good readability (and hence a proximity with the way agents express their preferences in natural language). In Section 3 we give a brief overview of several classes of logical languages for preference representation, and then we focus on two of these families of languages (both aiming at representing compactly utility functions): first, *weighted logics*, which express utility functions compactly by associating to each formula (representing an elementary goal) a positive weight representing the *reward* for satisfying this goal or the *penalty* for not satisfying it (together with a way of aggregating rewards and penalties stemming from different satisfied or violated goals); second, *distance-based logics*, where the violation of a goal is graded using a *distance* between worlds: given a goal (encoded as a propositional formula) G and given a distance d between worlds, then the closer a world w to a model of G, the better w (the optimal case being when w satisfies φ).

In [17] we showed that in some specific cases, a goal base written within the framework of distance-based logics can be translated into a goal base written within the framework of weighted logics and vice-versa.