Artificial Intelligence 175 (2011) 1153–1179

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Decision making with multiple objectives using GAI networks

C. Gonzales, P. Perny *, J.Ph. Dubus

LIP6 – Université Pierre et Marie Curie, case 169, 4 place jussieu, 75005 Paris, France

ARTICLE INFO

Article history:
Received 28 February 2009
Received in revised form 20 July 2010
Accepted 15 September 2010
Available online 2 December 2010

Keywords:
Graphical models
GAI decomposable utility
Preference representation
Multiobjective optimization
Multiagent decision making
Compromise search
Fairness

ABSTRACT

This paper deals with preference representation on combinatorial domains and preference-based recommendation in the context of multicriteria or multiagent decision making. The alternatives of the decision problem are seen as elements of a product set of attributes and preferences over solutions are represented by generalized additive decomposable (GAI) utility functions modeling individual preferences or criteria. Thanks to decomposability, utility vectors attached to solutions can be compiled into a graphical structure closely related to junction trees, the so-called GAI network. Using this structure, we present preference-based search algorithms for multicriteria or multiagent decision making. Although such models are often non-decomposable over attributes, we actually show that GAI networks are still useful to determine the most preferred alternatives provided preferences are compatible with Pareto dominance. We first present two algorithms for the determination of Pareto-optimal elements. Then the second of these algorithms is adapted so as to directly focus on the preferred solutions. We also provide results of numerical tests showing the practical efficiency of our procedures in various contexts such as compromise search and fair optimization in multicriteria or multiagent problems.

1. Introduction

The complexity of decision problems in organizations, the importance of the issues raised and the increasing need to explain or justify any decision has led decision makers to seek a scientific support in the preparation of their decisions. For many years, rational decision making was understood as solving a single-objective optimization problem, the optimal decision being implicitly defined as a feasible solution minimizing a cost function under some technical constraints. However, the practice of decision making in organizations has shown the limits of such formulations. First, there is some diversity and subjectivity in human preferences that requires distinguishing between the objective description of the alternatives of a choice problem and their value as perceived by individuals. In decision theory, alternatives are often seen as multiattribute items characterized by a tuple in a product set of attributes domains, the preferences of each individual being encoded by a utility function defined on the multiattribute space measuring the relative attractiveness of each tuple. Hence the objectives of individuals take the form of multiattribute utility functions to be maximized. Typically, in a multiagent decision problem, we have to deal with several such utility functions that must be optimized simultaneously. Since individual utilities are generally not commensurate, constructing an overall utility function gathering all relevant aspects is not always possible. Hence the problem does not reduce to a classical single-objective optimization task; we have to solve a multiobjective problem.

Moreover, even when there is a single decision maker, several points of views may be considered in the preference analysis, leading to the definition of several criteria. Rationality in decision making is generally not only a matter of costs reduction. In practice, other significant aspects that are not reducible to costs must be included in the analysis; the outcomes

* Corresponding author.
E-mail addresses: kaveh@river-valley.com (C. Gonzales), patrice.perny@lip6.fr (P. Perny), cvr@river-valley.com (J.Ph. Dubus).

0004-3702/$ – see front matter
doi:10.1016/j.artint.2010.11.020

of alternatives must be thought in a multidimensional space. This is the case in the elaboration of public policies where different aspects such as ecology and environment, education, health, security, public acceptability are considered in the evaluation process. This is also the case for individual decision of consumers. For example, when choosing a new car for a family, an individual will look at the cost, but will also consider several multiattribute utility functions concerning security in the car (brake system, airbags,...), velocity (speed, acceleration,...), space (boot size,...), environmental aspects (pollution) and aesthetics (color, shape, brand,...). All these observations have motivated the emergence of multicriteria methodologies for preference modeling and human decision support [1–4], an entire stream of research that steadily developed for forty years.

As for human decision making, automated decision making in complex environment requires optimization procedures involving multiple objectives. This is the case when computers are used for planning actions of autonomous agents or for organizing the workflow in production chains. Various other examples can be mentioned such as web search [5], *e*-commerce and resource allocation problems. In many of them, however, a decision is actually characterized by a combination of local decisions, thus providing the set of alternatives with a combinatorial structure. This explains the growing interest for multiobjective combinatorial optimization. Besides the explicit introduction of several possibly conflicting objectives in the evaluation process, the necessity of exploring large size solution spaces is an additional source of complexity. This has motivated the development in the AI community of preference representation languages aiming at simplifying preference handling and decision making on combinatorial domains.

As far as utility functions are concerned, the works on compact representation aim at exploiting preference independence among some attributes so as to decompose the utility of a tuple into a sum of smaller utility factors. Different decomposition models of utilities have been developed to model preferences. The most widely used assumes a special kind of independence among attributes called "mutual preferential independence". It ensures that preferences are representable by an additively decomposable utility [6,7]. Such decomposability makes both the elicitation process and the query optimizations very fast and simple. However, in practice, preferential independence may fail to hold as it rules out any interaction among attributes. Generalizations have thus been proposed in the literature to significantly increase the descriptive power of additive utilities. Among them, *multilinear utilities* [2] and GAI (generalized additive independence) decompositions [8,9] allow quite general interactions between attributes [7] while preserving some decomposability. The latter has been used to endow CP nets with utilities (UCP nets) both under uncertainty [10] and under certainty [11]. GAI decomposable utilities can be compiled into graphical structures closely related to junction trees, the so-called GAI networks. They can be exploited to perform classical optimization tasks (e.g. find a tuple with maximal utility) using a simple collect/distribute scheme essentially similar to that used in the Bayes net community or to variable elimination algorithms in CSP [12–15]. In order to extend the use of GAI nets to multiobjective optimization tasks, we investigate the potential of GAI models for representing and solving multiobjective optimization problems.

As soon as multiple criteria or utility functions are considered in the evaluation of a solution, the notion of optimality is not straightforward. Among the various optimality criteria, the concept of Pareto optimality or efficiency is the most widely used. A solution is said to be Pareto-optimal or efficient if it cannot be improved on one criterion without being depreciated on another one. Pareto optimality is natural because it does not require any information about the relative importance of criteria and can be used as a preliminary filter to circumscribe the set of reasonable solutions in multiobjective problems. However, in combinatorial optimization problems, the complete enumeration of Pareto-optimal solutions is often infeasible in practice [16–18]. For this reason, in many real applications, people facing such complexity resort to artificial simplifications of the problem, either by focusing on the most important criterion (as in route planning assistants), or by performing a prior linear aggregation of the criteria to get a single objective version of the problem, or by generating samples of good solutions using heuristics, which in any case does not provide formal guarantees on the quality of the solutions.

In this paper, we assume that each objective is represented by a GAI decomposable utility function defined on the multiattribute space describing items. In Section 2, after recalling basic definitions related to GAI nets, we show how they make it possible to represent vector-valued utility functions in a compact form, thus facilitating preference handling in multiobjective decision-making problems. In Section 3, we present two exact algorithms exploiting the structure of the GAI net for the determination of Pareto-optimal elements. In Section 4 we propose a refinement of the second algorithm aiming at focusing the search on specific compromise solutions within the Pareto set. We provide exact algorithms for preference-based search with various preference models. The potential of this approach is illustrated in the context of fair multiagent optimization or in the context of compromise search in multicriteria optimization. Finally, in Section 5, we present numerical tests showing the practical feasibility of the proposed approach on various instances of multiobjective combinatorial problems.

2. Multidimensional GAI nets

We assume that alternatives are characterized by n attributes $x_1, \dots, x_n$ taking their values in finite domains $X_1, \dots, X_n$ respectively. Hence alternatives can be seen as elements of the product set of these domains $\mathcal{X} = X_1 \times \cdots \times X_n$. In the sequel, $\mathbf{N} = \{1, \dots, n\}$ will denote the set of all the attributes' indices. By abuse of notation, for any set $\mathbf{Y} \subseteq \mathbf{N}$, $X_\mathbf{Y}$ will refer to the Cartesian product of the X_i, $i \in \mathbf{Y}$, i.e., $X_\mathbf{Y} = \prod_{i \in \mathbf{Y}} X_i$, and $x_\mathbf{Y}$ will refer to the projection of $x \in \mathcal{X}$ on $X_\mathbf{Y}$, that is, the tuple formed by the x_i, $i \in \mathbf{Y}$. We also consider a binary relation $\succsim$ over $\mathcal{X}$ (actually this is a weak order). Essentially, $x \succsim y$ means that x is at least as good as y. Symbol $\succ$ refers to the asymmetric part of $\succsim$ and $\sim$ to the symmetric one.