Available online at www.sciencedirect.com
ScienceDirect
Artificial Intelligence 172 (2008) 325–349
Artificial Intelligence
www.elsevier.com/locate/artint

Graphically structured value-function compilation

Ronen I. Brafman [a], Carmel Domshlak [b,*]

[a] *Department of Computer Science, Ben-Gurion University, Beer Sheva, Israel 84105*
[b] *Faculty of Industrial Engineering and Management, Technion—Israel Institute of Technology, Haifa, Israel 32000*

Received 24 January 2007; received in revised form 3 June 2007; accepted 11 July 2007

Available online 5 September 2007

Abstract

Classical work on eliciting and representing preferences over multi-attribute alternatives has attempted to recognize conditions under which value functions take on particularly simple and compact form, making their elicitation much easier. In this paper we consider preferences over discrete domains, and show that for a certain class of simple and intuitive qualitative preference statements, one can always generate compact value functions consistent with these statements. These value functions maintain the independence structure implicit in the original statements. For discrete domains, these representation theorems are much more general than previous results. However, we also show that it is not always possible to maintain this compact structure if we add explicit ordering constraints among the available outcomes.

Keywords: Reasoning about preferences; Qualitative decision theory; Multi-attribute decision making

1. Introduction

The spectrum of practical problems that require reasoning about preferences is extremely wide. In this paper we consider the problem of eliciting and reasoning about a user's *ordinal* preferences. We are motivated in part by the needs of large-scale, consumer product catalogs, an area that has received growing attention in the fields of the database systems and AI (e.g., see [1,5,6,9,13,26,29,31]).

Online catalogs of products and information grow continuously, and with them grows the number of lay users accessing these catalogs. While keyword search provides users with some means to access these catalogs, user needs in such shopping contexts are typically more complex than in web search. In particular, users have personal preferences regarding price, quality, features, etc., and these preferences can be rather complex. Therefore, it is natural to expect that systems supporting this search process will aim to allow users to state their actual preferences, and that reasoning about such preferences can improve the understanding of user needs.

Unfortunately, it appears that achieving *both* user-friendly, robust preference elicitation and efficient reasoning about the elicited information is not easy. The conflict between these two desiderata is reflected by the conflicting forms in which a user might be asked to provide her preferences. On the one hand, if the user provides us with a

* Corresponding author.
E-mail addresses: brafman@cs.bgu.ac.il (R.I. Brafman), dcarmel@ie.technion.ac.il (C. Domshlak).

0004-3702/$ – see front matter
doi:10.1016/j.artint.2007.07.002

numerical value function over the space of the products (henceforth referred to as *items*), ordering the catalog with respect to this function is easy. However, eliciting a quantitative description of preferences from the users is generally a long, involved and time-consuming process that is often unintuitive to users. Alternatively, we can consider allowing users to express their preferences using natural-language like *qualitative statements*, providing us with pieces of preference information like (i) "For a family car, I prefer white color to all other colors", or (ii) "This car is better for me than that car", or (iii) "This mini-van would be better in blue", or (iv) "I like ecologically friendly cars". This form of preference elicitation is considered to be more natural to users [18], and thus dealing with this form of preference information has received significant attention in the multi-disciplinary preference literature (to name just a few works, see [8,13,19,30,32]). Unfortunately, those preference expressions that can be reasoned about efficiently (at least for ordering a given set of items) are required to be "syntactically homogeneous", that is contain only statements in a certain specific form [8,14]. For instance, to the best of our knowledge, there is no known general class of preference expressions containing statements of both forms (i) and (ii) as above[1] that can be reasoned with efficiently.

Striving to enjoy the pros of both a qualitative input and a quantitative representation of user preferences, in this paper we consider *compiling* qualitative preference expressions into value functions consistent with the information carried by these expressions. The main contributions of this paper are as follows:

1. We provide a new *representation theory* for generalized additive value functions [2,15], and specify conditions under which there exists a particular factored value function consistent with (what is known about) the user's preference relation. Our representation theorems show that preference orders induced over the item space by certain sets of qualitative statements of preference and importance can always be consistently captured by a compact generalized additive value function. In particular, our results extend the classical representation theorems for additive value functions over discrete variables [21]. As the conditions we require are much weaker than those required for an additive representation, we are able to capture a significantly wider spectrum of sets of natural preference statements, namely those representable by the TCP-net model [10,12].
2. We show how our representation theory can be utilized in a *computationally efficient* methodology for eliciting and reasoning about ordinal preferences of the users. In this methodology, the user provides a set of qualitative preference statements, and these statements are used to efficiently generate a compact value function whose structure is based on the qualitative information supplied by the user. The key part is that the existence of such a compact value function, its consistency with the preference statements of the user, and efficiency of its generation are guaranteed by our representation theory.
3. In many applications, it is desirable to allow the users expressing not only structured preference information, but also direct rankings between pairs of concrete items (e.g., see [5,20,26]). We consider the computational consequences of supporting both general statements of preference and such pair-wise item rankings. On the positive side, we show that such an extension can be straightforwardly supported in our methodology while preserving its soundness and efficiency. On the negative side, however, we formally show that completeness of structured value-function compilation is extremely sensitive to adding such item-level rankings. Specifically, we show that completeness of value-function compilation cannot be guaranteed even if the amount of such pair-wise item rankings is minimal, and that this impossibility result holds for most languages of generalizing preference statements.

The rest of the paper is organized as follows. In Section 2 we provide some essential background on qualitative preference statements targeted in this work, the TCP-model for modeling sets of such statements, and value functions. Section 3 is devoted to the value-function compilation of three progressively more complicated classes of TCP-nets. For clarity of presentation, the longer proofs are given in Appendix A. In Section 4 we consider extending structured preference information with pairwise comparisons between completely specified alternatives, provide an impossibility theorem on value-function compilation of such mixed sets of statements, and generalize this result to a general impossibility theorem. We summarize and list some open problems in Section 5.

[1] Later we define these forms of statements in a formal manner.