Reasoning About Multi-Attribute Preferences

(Extended Abstract)

Koen Hindriks Catholijn Jonker Wietske Visser
Man Machine Interaction Group, Delft University of Technology
Mekelweg 4, 2628 CD Delft, The Netherlands
K.V.Hindriks@tudelft.nl, C.M.Jonker@tudelft.nl, Wietske.Visser@tudelft.nl

1. INTRODUCTION

In our aim to develop a negotiation support system, we are faced with the need to express a user's preferences. The offers that are exchanged usually consist of multiple attributes. An agreement can only be reached if the preferences from both parties over these attributes and the complete offers are taken into account. The acceptablilty of an offer is defined in terms of the value that a negotiator attaches to it.

Unfortunately, eliciting and representing a user's preferences is not unproblematic. Existing negotiation support systems are based almost exclusively on quantitative models of preferences. These kinds of models are based on utilities; a utility function determines for each outcome a numerical value of desirability. However, it is difficult to elicit such models from users, since the required numerical data are not easily available. Humans generally express their preferences in a more qualitative way. We say we like something more than something else, but it is certainly strange to express liking something exactly twice as much as an alternative. Therefore, qualitative preference models would provide a better correspondence with human cognitive representations. This would allow a human user to interact more naturally with a machine negotiating on his behalf.

Object preferences can be quite complex. For example, it is quite natural to say that you prefer one house over another because it is bigger and generally you prefer larger houses over smaller ones. So it is a natural approach to derive object preferences from general property or attribute preferences. Besides, the number of objects is combinatorially large in the number of properties, so it soon becomes infeasible to specify object preferences directly.

How can preferences over objects be expressed? One approach to obtain preferences about objects is to start with a set of properties of these objects and derive preferences from a ranking of these properties, where the ranking indicates the relative importance or priority of each of these properties. This approach to obtain preferences is typical in multi-attribute decision theory, see e.g. Keeney and Raiffa [8]. Multi-attribute decision theory provides a quantitative theory that derives object preferences from utility values assigned to outcomes which are derived from numeric weights associated with properties or attributes of objects. As it is difficult to obtain such quantitative utility values and weights, however, several qualitative approaches have been proposed instead, see e.g. [1, 3, 4, 5, 9]. Doyle and Thomason [6] provide an overview of central questions of interest in qualitative decision theory.

Cite as: Reasoning About Multi-Attribute Preferences, (Extended Abstract), Koen Hindriks, Catholijn Jonker, Wietske Visser, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. 1147–1148

Key to a logic of multi-attribute preferences is the representation of property rankings. Encodings of property rankings have been explored in Coste-Marquis *et al.* [5] where they are called goal bases, and in Brewka [4] where they are called ranked knowledge bases. Coste-Marquis *et al.* [5] and Brewka [4] moreover discuss various options, or strategies, for deriving object preferences from a property ranking. The preferences orderings thus obtained are not expressed in a logic, however. Brewka *et al.* [3] propose a non-monotonic logic called qualitative choice logic to reason about multi-attribute preferences. An alternative approach towards a logic of multi-attribute preferences is presented in Liu [9] where property rankings called priority sequences are encoded in first-order logic. Both approaches are based on one particular strategy, namely lexicographic ordering, to obtain object preferences from a property ranking.

2. PREFERENCE ORDERINGS

We propose a generic logic of qualitative multi-attribute preferences, in which property rankings and associated strategies for deriving object preferences from such rankings can be defined. Various such strategies can be found in the literature. Coste-Marquis *et al.* [5] describe three frequent orderings based on prioritized goals: best-out, discrimin and leximin ordering. Brewka [4] defines a preference language in which different basic preference orderings can be combined and identifies four 'fundamental strategies' for deriving preferences from what he calls a ranked knowledge base: $\top$, κ, $\subseteq$ and #. Three of Brewka's defintions are equivalent to the definitions of Coste-Marquis *et al.*

The two orderings $\subseteq$ and # first consider the most important property. If some object has that property and another does not, then the first is preferred over the second. If two objects both have the property or if neither of them has it, the next property is considered. The $\subseteq$ and # orderings only differ if multiple properties are equally important. The $\top$ ordering looks at the highest ranked or most important property that is satisfied. If that property of one object is ranked higher than that of another object, then the first object is preferred over the second. If those properties are equally ranked, then both objects are equally preferred. The κ ordering looks at the most important property that is not satisfied. If that property of one object is less important than the property of another object, then the first object is preferred over the second. If those properties are equally important, then both objects are equally preferred.

Expressing these preference ordering strategies and the property rankings from which the orderings are derived in a logic allows us to reason with and about qualitative multi-attribute preferences.

3. MULTI-ATTRIBUTE PREFERENCE LOGIC

The multi-attribute preference logic we propose is a modal logic that extends binary preference logic as presented in Girard [7] with names for objects and a modal operator to characterize properties or attributes. This language extension allows us to talk about properties, objects and associated preferences explicitly.

The basic concepts in the semantics for multi-attribute preference logic are objects and properties those objects may have. Properties are naturally represented by sets of worlds in modal semantics, which we call clusters. As we want to use properties to classify the ranking of objects, properties are ordered in correspondence with their relative importance; such an order is called a property ranking.

Objects are identified with particular sets of worlds. The idea is that we can derive the properties of an object from the worlds which define the object. To ensure that objects are coherent, that is, have a uniquely defined set of properties, the worlds that define the object need to be copies of each other. A world is a copy of another world if it assigns the same truth values to propositional atoms. The general idea is that the worlds that constitute an object act as representatives for that object in a certain property cluster. So for each object, there is a world for each property in the property ranking that the object has.

One of the binary preferences in [7] is expressed as $\phi <_{\forall\forall} \psi$, which expresses that any state where ψ is true is strictly better than any state where ϕ is true. If we use i and j for object names, then $i <_{\forall\forall} j$ expresses that object j is preferred over i. The preference expressed in this way is a very strong kind of preference, however. It requires that all of object j's relevant properties are considered more important than objects i's properties. The motivation to introduce multi-attribute preference logic is to provide a logic that enables the specification of principles that allow to derive preferences over objects from their properties in a weaker sense. In most qualitative preference orderings, for example, it is possible that object i has at least one property that is considered more important than a property that object j has but object j is still preferred over object i. Moreover, the logic should provide the means to derive a preference of one object over another from a specification of the properties of these objects and a priority sequence associated with these properties, or a property ranking.

As said above, Coste-Marquis *et al.* [5] and Brewka [4] have defined various strategies to obtain object preferences from a property ranking. All these preference orderings can be defined in multi-attribute preference logic in such a way that they are equivalent to those of Brewka. A ranked knowledge base, which is Brewka's version of a property ranking, can be translated into multi-attibute preference logic in such a way that every multi-attribute preference model that is a model of the translation yields the same preference ordering as the original ranked knowledge base.

The advantage of defining preference orderings in a logic instead of providing set-theoretical definitions is that it formalizes the reasoning about object preferences. From a practical point of view, the logic allows to provide rigorous formal proofs for object preferences derived from property rankings. From a theoretical point of view, it provides the tools to reason about preference orderings and allows, for example, to prove that whenever an object is preferred over another by the $\top$ strategy it also is preferred by the # strategy.

4. FUTURE WORK

An interesting issue to explore further is how to express dependencies between attributes. Existing work on CP-nets [2], which enable to define conditional preferences, might provide a direction for further investigation of this issue.

There are more challenges in preference representation. Preferences may not be complete. Users may lack some of the information that is relevant for the comparison, or they may have inconsistent or uncertain beliefs, or have undetermined or incoherent preferences. To resolve this issue, either the missing information must be acquired, or we must deal with incompleteness and derive preferences nonetheless in a nonmonotonic approach. Nonmonotonicity also arises in case of preference change. Preferences may change due to various factors, such as experience, reasoning, action, communication with others, and changing goals and interests. Logics for defeasible argumentation [10] might provide an interesting approach to tackle the issue of nonmonotonicity. Argument-based reasoning is a natural way to reason, and it might be better equipped to deal with inconsistent, incomplete or changing beliefs and preferences. Moreover, it provides a computational approach to derive preferences.

A general downside of qualitative approaches to preferences is that it is hard to deal with pay-offs between attributes. Hybrid approaches that combine quantitative and qualitative preferences are a very interesting angle to explore further. This would also be helpful in negotiation support, where strategies might depend on the size of a concession made by the opponent.

Eventually, our aim is to integrate an expressive preference language into a larger negotiation framework. This framework will also contain associated strategies for negotiation and will be the core of a negotiation support system.

5. REFERENCES

[1] C. Boutilier. Toward a logic for qualitative decision theory. In *Proc. KR 1994*, pages 75–86, 1994.

[2] C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artificial Intelligence Research*, 21:135–191, 2004.

[3] G. Brewka. A rank based description language for qualitative preferences. In *Proc. ECAI 2004*, 2004.

[4] G. Brewka, S. Benferhat, and D. Le Berre. Qualitative choice logic. *Artificial Intelligence*, 157(1-2):203–237, 2004.

[5] S. Coste-Marquis, J. Lang, P. Liberatore, and P. Marquis. Expressive power and succinctness of propositional languages for preference representation. In *Proc. KR 2004*, pages 203–212, 2004.

[6] J. Doyle and R. H. Thomason. Background to qualitative decision theory. *AI Magazine*, 20(2):55–68, 1999.

[7] P. Girard. *Modal Logic for Belief and Preference Change*. PhD thesis, Universiteit van Amsterdam, 2008.

[8] R. L. Keeney and H. Raiffa. *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge University Press, 1993.

[9] F. Liu. *Changing for the Better: Preference Dynamics and Agent Diversity*. PhD thesis, Universiteit van Amsterdam, 2008.

[10] H. Prakken and G. A. W. Vreeswijk. Logics for defeasible argumentation. In D. Gabbay and F. Günthner, editors, *Handbook of Philosophical Logic*, volume 4, pages 219–318. Kluwer Academic Publishers, 2002.