

# User-Centric Preference-Based Decision Making

## (Extended Abstract)

Ingrid Nunes  
PUC-Rio, LES  
King's College London  
ionunes@inf.puc-rio.br

Simon Miles, Michael Luck  
King's College London  
London, United Kingdom  
{simon.miles,michael.luck}  
@kcl.ac.uk

Carlos J.P. de Lucena  
PUC-Rio, LES  
Rio de Janeiro, Brazil  
lucena@inf.puc-rio.br

### ABSTRACT

The automation of user tasks by agents may involve decision making that must take into account user preferences. This paper introduces a decision making technique that reasons about preferences and priorities expressed in a high-level language in order to choose an option from the set of those available. Our technique includes principles from psychology, concerning the way in which humans make decisions. Our preference language is informed by a user study on preference expression, which is also used to evaluate our approach by comparing our results with those provided by a human expert. The evaluation indicates that our technique makes choices on behalf of the user with as good quality as made by the expert.

### Categories and Subject Descriptors

H.4.2 [Information Systems Applications]: Types of Systems—Decision support; I.2.11 [Computing Methodologies]: Distributed Artificial Intelligence—Intelligent agents

### General Terms

Algorithms

### Keywords

Agent Reasoning, Preferences, Decision Making, Autonomous Agents

## 1. INTRODUCTION

The automation of user tasks by agents may involve decision making that must take into account user preferences. Our vision is for agents to make decisions for users so that their choices match those of users themselves, given adequate time and knowledge. People do not act in isolation, and agents acting on their behalf should not do so either. Where the option chosen for one user may affect that of another (e.g., in deciding which hotel to stay at, we both prefer to stay at the same hotel), agents need to coordinate their actions. Such coordination between users reflects just one among the many interacting preferences that agents may need to consider. We argue that, by reflecting how users themselves decide, there is a *rationale* for choices that is *convincing* to users.

This paper describes the first step towards this vision. Before we can have decisions appropriate to multiple users, we must first have agent reasoning appropriate to a single user. However, many

**Appears in:** *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

Copyright © 2012, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

different approaches have been proposed for reasoning about preferences, but they address a restricted set of preference types, and therefore are not able to process preferences provided by users in many realistic scenarios. We propose a novel approach for reasoning about preferences. Specifically, the contributions of this work are (i) a high-level preference language, informed by user preferences in natural language; and (ii) an automated decision-making technique based on preferences and available options, and exploiting psychological research into the way in which humans make choices.

Our decision-maker takes as input a set of options over which a choice is made, and a set of preferences expressed in the language introduced below. It processes the preferences to select one option, in such a way that the choice, and the decision not to choose alternatives, can be justified by the preferences. The output is a partially ordered set, organised in four different levels: (i) the *chosen option*; (ii) *acceptable options* that are *close* to the chosen option, but not chosen; (iii) *eliminated options*, discarded because of a hard constraint; and (iv) *dominated options*. We apply heuristics used by humans, specifically the principles of *trade-off contrast* and *extremeness aversion* [3], so that decisions more closely mirror user choices if users are provided with sufficient time and knowledge. In outline, the steps of our technique are as follows.

**Pre-processing.** Options are analysed to extract the essential data. This includes how well option attributes meet the preferences, and how options compare with regard to individual attributes.

**Explication.** Some preferences include important *implicit* information, in addition to their literal meaning, and we extract it.

**Elimination.** Next, we eliminate options that do not meet strict constraints, or that are dominated in every regard by other options.

**Selection.** Finally, we make the choice itself. As the remaining options have both costs and benefits, we need to take account of all preferences that lead to a decision, such as the relative importance of attributes, plus heuristics, including the principles of *trade-off contrast* and *extremeness aversion* adopted by humans [3].

## 2. PREFERENCE LANGUAGE

Humans express preferences in many ways, and we wish to provide them with this natural expressivity when delegating decisions to a software agent. We propose a preference language derived from an existing user study on choosing a laptop, based on around 200 preference specifications [1]. Our language defines seven types of preference: *constraints* specify values that attributes must (not) have; *goals* specify which attributes should be minimised or maximised; *orders* specify preferences over attributes; *qualifying* preferences state how much an attribute value is wanted or needed; *rating* preferences specify which values are best or worst; *indifferences* specify the absence of preference between two attribute

values; and *don't care* preferences specify irrelevant attributes. In addition, preferences may apply only conditionally, where the *condition* is expressed in terms of attribute values, and *priorities* can be expressed either between attributes or preferences, so that an attribute or preference is given more weight in decision-making.

### 3. PRE-PROCESSING

Preferences can be *monadic* or *dyadic*, where the former evaluate a single referent, e.g. an apartment less than 2.5km away from the university is preferred, and the latter indicate a relation between two referents, e.g. lower price is better. First, we pre-process the options with regard to monadic and dyadic statements, thus building two models for use in later steps.

Performatives such as *need*, *require*, and *love* are widely adopted by users to express preferences over attributes, and so are included in our language. Similarly, users may rate preferences from best to worst. The rates and performatives used in monadic preferences are captured by a *Preference Satisfaction Model (PSM)*, which consists of a table indicating how options satisfy monadic preferences in terms of each attribute. The Options-Attribute Preference Model (OAPM) is a table that captures comparisons between two options, for individual attributes, showing which is better, or that no conclusion can be drawn from the provided preferences. The OAPM is based on preferences not used in the PSM, together with the PSM itself, processed separately in a specific order, and establishing a precedence: (i) order and indifference; (ii) goals; and (iii) PSM.

### 4. EXPLICATION AND ELIMINATION

Preferences always provide a literal meaning, but can also bring additional information to derive new preferences, referred to as *implicit preferences*. These never override information of explicitly provided preferences, but enable determination of whether an option is preferred to another with respect to a certain attribute, when this could not otherwise be concluded. We update the OAPM by considering these implicit preferences, such as considering that a higher value of an attribute (maximisation goal) is better than a lower one, if there is a preference that establishes a lower bound for this attribute, and both options satisfy this preference.

A typical approach adopted by users in making a choice is the stepwise elimination of options until there remains a set of acceptable options, ideally containing only one element, as in *elimination by aspects* [4]. In the elimination step, we discard two types of options: (i) dominated options; and (ii) options that do not satisfy hard constraints. The OAPM and the PSM are used to identify these options, respectively.

### 5. SELECTION

After *elimination*, we must choose an option from the acceptable set, i.e. available options without those eliminated. Humans commonly make use of heuristics [2], that demand different amounts of effort, typically choosing them by matching the effort required to the importance of the decision. Our approach does not aim to reproduce this behaviour, which relies on human decisions on investment of effort, but instead seeks to understand *how* users resolve trade-offs, regardless of the effort made.

We begin the process of choosing an option by evaluating each pair of options and assessing their costs and benefits. First, we analyse the benefits of option  $o_1$  compared to option  $o_2$  for each attribute, and do the same for  $o_2$  compared to  $o_1$ . Benefits are captured by a real value from 0 to 1, indicating how much better one option is than another, with respect to one attribute. If the OAPM indicates that  $o_1$  is not better than  $o_2$  for an attribute *att*, then the

benefit is 0, otherwise, to compute this benefit, we use the preference used to set the OAPM value. Having considered attributes in isolation, we now examine overall option benefits, via the priorities provided. Based on priorities, we build an attribute partial order, associating one level with each attribute. A function is adopted to generate attribute weights, and we calculate the overall benefits from  $o_1$  with respect to  $o_2$  using a weighted sum. It is important to highlight that benefits are obtained solely from high-level preferences, without requiring further interactions with the user.

If there are no dominated options in the set of acceptable options then, for any two options, one is better for some attributes and the same applies for the other. As a consequence, a trade-off must be resolved to choose one of the two options. According to Simonson and Tversky [3], people not only consider the two options being compared and their costs and benefits, but also the cost and benefit relationship (ratio), which is positioned in relation to this ratio between other options. This is referred to as *trade-off contrast*. In addition, humans also consider how *extreme* options are. Extreme options have a large improvement for some attributes, e.g. quality, and a high penalty for others, e.g. price. In general, humans avoid extreme options [3], referred to as *extremeness aversion*. We therefore incorporate two new factors in the process of choosing an option, based on a function that shows the trade-off between two options and how extreme they are.

We have analysed three aspects of options: benefits, trade-off relative to available options, and extremeness. The last two aspects are also seen as benefits (or costs): if the trade-off between two options is better according to the average of the trade-off between every other pair of options, it is also a benefit, and the less extreme option has a benefit in comparison to the more extreme. The final value of an option with respect to another is thus a weighted sum of these benefits. Based on the  $v$  function, we identify the chosen option as better than or equal to every other option. If different options have the same value with respect to another, and they are better than every other option, we randomly choose one of them.

### 6. CONCLUSIONS

Intelligent agents provided with mechanisms that enable them to reason about preferences and make choices on behalf of users are a promising solution for reducing user effort in the automation of tasks. In this paper, we propose an automated decision making technique, which chooses an option from the set of those available based on preferences and priorities expressed in a high-level preference language. We improve decision-making by incorporating user-centric principles (trade-off contrast and extremeness aversion) that are not explicitly expressed as preferences. Based on an empirical evaluation, we can conclude that our technique makes choices as good as those of a (human) domain expert.

### 7. REFERENCES

- [1] I. Nunes, S. Barbosa, and C. Lucena. Understanding how users express preferences: a user study. Technical Report CS-2010-19, University of Waterloo, Canada, 2010.
- [2] J. W. Payne, J. R. Bettman, and E. J. Johnson. Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14(3):534–552, 1988.
- [3] I. Simonson and A. Tversky. Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research*, 29(3):281–295, 1992.
- [4] A. Tversky. Elimination by aspects: a theory of choice. *Psychological Review*, 79(4):281–299, 1972.