Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence

Advice Provision in Multiple Prospect Selection Problems

Amos Azaria and Sarit Kraus
Department of Computer Science
Bar Ilan University Ramat Gan, Israel
{azariaa1,sarit}@cs.biu.ac.il

Abstract

When humans face a broad spectrum of topics, where each topic consists of several options, they usually make a decision on each topic separately. Usually, a person will perform better by making a global decision, however, taking all consequences into account is extremely difficult. We present a novel computational method for advice-generation in an environment where people need to decide among multiple selection problems. This method is based on the prospect theory and uses machine learning techniques. We graphically present this advice to the users and compare it with an advice which encourages the users to always select the option with a higher expected outcome. We show that our method outperforms the expected outcome approach in terms of user happiness and satisfaction.

Introduction

It is hard to overestimate the importance of decision-making in life. People make decisions on a daily basis; some are trivial such as whether or not to eat ice-cream, while others are important such as which apartment to buy. "Choice Bracketing", termed by Read et al. (1999), designates the grouping of individual choices together into sets. "Broadly Bracketing" indicates that the decision-maker takes all choices into account when making his decision, while "Narrow Bracketing" indicates that the decision-maker isolates each choice from all other choices. Of course, in order to bracket several decisions, they must all be converted to a single scale (such as money), where a decision that turns out to be a bad choice for one topic may be balanced out by another decision which turns out to be a good choice for a different topic.

It has been shown that people tend to use narrow bracketing and usually treat each decision as if in isolation from all other decisions, which in many cases results in a poor choice (Read, Loewenstein, and Rabin 1999).

We tackle people's tendency to narrow bracketing using an environment where people need to decide among several independent selection problems, whether they prefer a guaranteed outcome or a higher, but uncertain, outcome. We build an agent which first learns peoples' preferences and generates a general human model. The agent then searches the space of all possible combinations that can be chosen by the user and, based on the human model, advises the user which combination would be most beneficial for him.

We would like to emphasize, that although we focus on advice provision, the agent may utterly replace the decision-maker.

The main contributions of the paper are the following: first, we present the algorithmic challenge that is associated with the human bracketing problem. Second, we developed a multi-stage procedure for providing humans with a combination of choices in multiple prospect selections aiming at maximizing human satisfaction. The innovation in the preference elicitation of the multi-stage procedure stems from the use of an off-line learning group. New users do not need to provide any additional information. Furthermore, the group elicitation procedure is done using simple problems in which humans can provide their real preferences and address the inability of people to recognize their preferences in complex decision problems. At last, we provide extensive experiments that show the success of our proposed method in terms of human satisfaction.

In recent work (2012) we propose a method for giving advice to users in an environment where users have several options from which to choose where the decisions are made in sequential order. However, in that work we assume that the users and the agent have different goals and utility functions, while in this paper the agent's goal is solely to help the user and thus they both have identical utility functions.

Multiple Prospect Selections Advice Provision

A *prospect selection problem*, is a problem where a player needs to choose between a guaranteed outcome of x and a probability of p to win an outcome y. An important point is that a fully rational player (i.e. one who maximizes expected monetary outcome) would always choose the uncertain outcome if $p \cdot y \geq x$, and the guaranteed outcome otherwise. However, as shown by Tversky and Kahneman (1992), people do not act fully rationally not because they cannot do so. Kahneman and Tversky show that people have a subjective representation of probabilities and do not interpret probabilities fully rationally, but rather use their own decision weights when deciding whether to reject or accept a gamble. We consider *multiple prospect selection problems* where a player faces k (different) prospect selection problems. We build an agent that advises a human player in which of the selection problems to choose the guaranteed outcome and in which to choose the prospect. We measure the performance of the agent in terms of human satisfaction.

We now present the prospect Selection problem Advice provider for Multiple Problems agent (SAMP). To build this agent we first use machine learning to elicit decision weights from given data. Based on these decision weights we build a general human model that can calculate the utility in human eyes for any combination of choices in a multiple prospect selection problem. Upon demand, SAMP searches for the best combination using this model and presents it to the user.

The first step in building SAMP, is eliciting human decision weights. We collect data from subjects who were asked to choose between a guaranteed outcome and a simple prospect. Note that this is a very simple choice and we therefore assume that humans can provide their real preferences. Using the data we build a logistic regression classifier to elicit the decision weights.

The main challenge which remains is to assess the value of multiple (k) prospects. If we were simply to sum the value of all prospects individually, we would fail by using the exact concept which we are trying to overcome, i.e. narrow bracketing. The first step in assessing the value of multiple prospects is calculating the final probability for each of the possible outcomes. Once we obtain the probabilities $\vec{p} = p^1, p^2 \dots p^n$ ($\sum_{i=1}^{n} p^i = 1$) and outcomes $\vec{y} = y^1, y^2 \dots y^n$ ($y^1 < \dots < y^n$), based upon the Cumulative Prospect Theory (Tversky and Kahneman 1992), we assess the value for the user by: $u(\vec{p}, \vec{y}) = \sum_{i=1}^{n} d\big(\sum_{j \geq i} p^j\big) y^i - d\big(\sum_{j > i} p^j\big) y^i$. Given a Multiple Prospect Selection Problem, each combination of choices yields different vectors $\vec{p}$ and $\vec{y}$ and therefore using the model above, each combination of choices yields a different user value ($u(\vec{p}, \vec{y})$). SAMP searches for a combination of choices which maximizes this value. We compare the performance of SAMP to the performance of a *rational agent* which assumes human full rationality (i.e. assumes that people want to maximize their expected monetary value). Performance is measured in terms of human satisfaction.

Experimental Design

We recruited 52 participants for SAMP's learning phase and 202 participants for evaluating both SAMP and the fully rational agent. In the evaluation phase, each subject received a single advice, either from SAMP or from the fully rational agent. In the learning phase the subjects did not receive any advice. Each subject participated only once. We set $k = 5$, i.e. the subjects had to make their choice regarding five prospect selection problems. Prior to receiving the actual advice, the subjects were told that the advice is provided by a third party agent which is trying to help them. The subjects were shown the consequences of following the advice using a pie chart, which indicated the actual probability of winning each possible outcome.

Figure 1 presents an example for a visualization of the agent's advice. In this example, if the user follows the exact advice given by the agent, he is guaranteed a win of at least 8 cents, however, he might win upto 89 cents. The interpretation of the pie chart was explained to the subjects, and their comments indicated that they clearly understood this interpretation. The subjects were asked the following three questions: 1. Are you happy with the final result? 2. Are you happy with the decisions you made? 3. How good was the advice given to you by the system? The subjects answered these questions on a 1 to 5 scale.

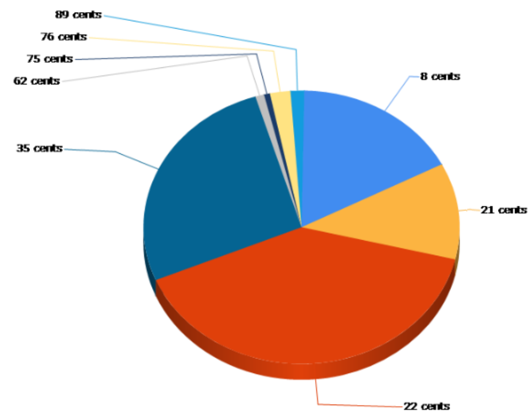

Figure 1: A visualization of the agent's advice

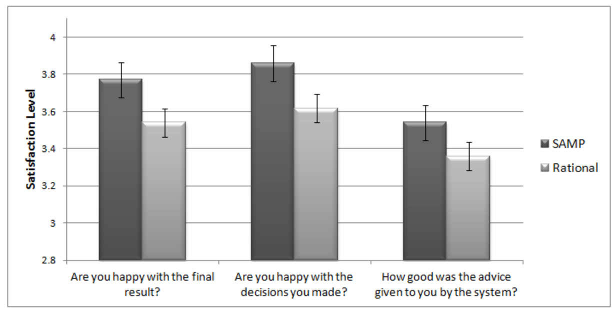

Figure 2: Comparison between user satisfaction levels among subjects who received SAMP's advice and those who received the fully rational agent's advice

Figure 2 presents the final satisfaction levels for users who received SAMP's advice and users who received the fully rational advice. The higher the satisfaction level the better the result. As can be seen in the figure, in all three questions SAMP outperformed the rational agent. ANOVA test indicates that these results differ significantly ($p < 0.05$).

References

Azaria, A.; Rabinovich, Z.; Kraus, S.; Goldman, C. V.; and Gal, Y. 2012. Strategic advice provision in repeated human-agent interactions. In *AAAI*.

Read, D.; Loewenstein, G.; and Rabin, M. 1999. Choice bracketing. *Journal of Risk and Uncertainty* 19(1):171–197.

Tversky, A., and Kahneman, D. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4):297–323.