

A System for Advice Provision in Multiple Prospect Selection Problems

Amos Azaria¹ and Sarit Kraus^{1,2} and Ariella Richardson³

¹ Department of Computer Science, Bar-Ilan University, Ramat Gan 52900, Israel

² Institute for Advanced Computer Studies University of Maryland, MD 20742

³ Jerusalem College of Technology, Jerusalem 91160, Israel
{azariaa1,sarit}@cs.biu.ac.il, richards@jct.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 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 and satisfaction.

1. 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. [7], 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 is a bad choice for one topic may be balanced out by another decision which is 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 [10, 7]. Gneezy and Potters [2] have shown that, in investment games, people who are forced into broadly bracketing by viewing their revenue less often become less risk averse and therefore increase their average revenue. Unfortunately, these studies did not test users' satisfaction, only their average total revenue.

In this paper 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 in terms of satisfaction. Evaluating the performance of a system using user satisfaction is a common method in recommender systems [4]. A user is more likely to buy, use and recommend a system with high user satisfaction than without.

Many real life situations resemble our problem. The most obvious example is when building an investment portfolio, where some stocks have a higher risk but also offer an opportunity to receive higher interest and others a lower risk and a lower interest level. Most people combine different stocks and bonds, combining different levels of risk. A similar example might be a manager who wants to market several products, where each product has its own probability for success and failure and its estimated revenue. Another example might be a GPS guiding a driver who can choose between a longer road with a low probability of heavy traffic and a shorter road with a high probability for heavy traffic. Considering the full route may result in greater satisfaction than making a separate decision for each and every junction. We emphasize, that although we focus on advice provision, the agent may utterly replace the decision-maker. For instance, if the GPS is embedded into an autonomous vehicle, it may determine the route on its own.

Our main contributions are: first, we present the algorithmic challenge associated with the human bracketing problem. Second, we developed a multistage 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 multistage procedure stems from the use of an off-line learning group. At last, extensive experiments show the success of our proposed method in terms of human satisfaction.

2. RELATED WORK

In a recent work, Azaria et al. [1] 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 their work they assume that the users and the agent have different goals and utility functions, while in our work the agent's goal is solely to help the user and thus they both have identical utility functions. Sarne et al. [9] attempt to facilitate people's decision-making process by modifying the presentation of the problem in an economic search environment. However, in their work they assume that people want to maximize their expected monetary value. We relax this assumption and assume that people have a non-monetary utility function which they try to maximize. In this paper we base this utility function on the Prospect Theory.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

The Prospect Theory was presented by Kahneman and Tversky in [6], and later refined to the Cumulative Prospect Theory in [11]. The Prospect Theory is based on three principles. The first is that people do not take into account their total wealth when accepting or rejecting an uncertain opportunity but rather use their current wealth as a baseline and are happy to win an amount and upset if they lose. The second principle is loss aversion, and the third is 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. Tversky and Kahneman elicited these weights by sequentially asking subjects to choose between a specific lottery and many different guaranteed outcomes. However, these decision weights depend on peoples personality, their wealth, culture and the scope of the payoff in question. Several studies try to estimate parameters for the Prospect Theory [3, 8], however, most studies try to maximize the likelihood of the results obtained by each subject individually. This approach could not be used in our work since we build a model using a group of users and use this model on new users (for whom we have little or no data). We cannot build a model for each and every user - it is impractical and we don't have enough data for that. We build a human model by learning the decision weights which best suit our population and scope of payoff based on machine learning.

3. MULTIPLE PROSPECT SELECTIONS

The term *prospect* comes from Kahneman and Tversky, who refer to a lottery where a player has a chance to win (or lose) a certain amount as a prospect. We only use *simple prospects*, where a player can gain a certain outcome with some probability, and zero otherwise. A *prospect selection problem* $s = \langle x; p, y \rangle$, 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 (we use the subscripts s_x, s_p and s_y respectively). We use $c \in \{g, u\}$ to denote the user's choice, where g denotes a choice of the guaranteed value (x), and u denotes the uncertain outcome (the prospect: y with probability p). 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 [11], people do not act fully rationally not because they cannot do so, but because they use different decision weights.

In our work we consider *multiple prospect selection problems* where a player faces k (different) prospect selection problems. As mentioned in the introduction, people fail to broaden their bracketing and treat each selection problem as if it is isolated from the other selection problems. Therefore, people use their original decision weights, which would imply a lower satisfaction rate for multiple selection problems. Our goal is to 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 denote this advice by $a \in \{g, u\}$. We measure the performance of the agent in terms of human satisfaction from the final result, human satisfaction from the decisions he made, human satisfaction from the advice and the fraction of selection problems where the human followed the advice. We intentionally do not measure success in terms of the raw final stake, since, as mentioned above, people care about how much they have only in context, that is, in comparison to how much they could have had (or lost), and what risks they took or avoided in order to achieve that outcome.

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 which are most appropriate for a given domain and demographic group. Based on

these decision weights we build a general human model that can calculate the weighted utility for any combination of choices in a multiple prospect selection problem. On demand, SAMP uses the model to search for the best combination and presents it to the user.

3.1 Decision Weights Elicitation

The first step in building SAMP, is eliciting human decision weights. Since these decision weights will be used to build a *general* human model, we collect data using subjects from a similar culture and using a similar scope of outcome and probabilities. However, we do not collect personal parameters regarding users' preferences (such as risk attitude). This data is collected from subjects who were asked to choose between a guaranteed outcome and a simple prospect (as described earlier). Since this is a very simple choice, we assume that humans can provide their real preferences.

Eliciting these decision weights is not a classification problem, however, we use a classifier, in order to extract its model and then using its model along with mathematical manipulations we obtain a decision weight function. Using the data we build a logistic regression classifier. We feed the classifier for each prospect selection problem with the probability of winning in the prospect, the ratio between the prospect outcome and the guaranteed outcome, and all quadratic combinations of the two, using the the feature vector $v = \{s_p, \frac{s_y}{s_x}, s_p^2, (\frac{s_y}{s_x})^2, s_p \cdot \frac{s_y}{s_x}\}$. The classifier needs to classify the data based on the subjects' choices, and is trained on c .

Given a new prospect selection problem feature vector the classifier will output a number between 0 and 1 which, if greater than 0.5, indicates that the user is more likely to choose one option and if smaller than 0.5 indicates that the user is more likely to choose the other. However, since we are interested in learning the decision weights, we would like to know when people are *indifferent* between choosing the guaranteed and the uncertain outcome. We therefore are interested in the cases where the classifier will output exactly 0.5. Since we use a logistic regression classifier, we obtain $\frac{1}{1+e^{-(w^T \cdot v)}} = 0.5$ where w is the vector of weights obtained from the classifier, and v is the feature vector. This implies $(w^T \cdot v) = 0$.

Since we are interested in finding the decision weights for a given probability, we denote: $z = \frac{s_y}{s_x}$ and solve for z . Writing both the feature vector and the weight vector explicitly we obtain:

$$w_0 + w_1 \cdot s_p + w_2 \cdot z + w_3 \cdot s_p^2 + w_4 \cdot z^2 + w_5 \cdot s_p \cdot z = 0$$

and after solving we obtain:

$$z(s_p) = \frac{-w_5 - w_2 \pm \sqrt{(w_5 s_p + w_2)^2 - 4w_4(w_3 s_p^2 + w_1 s_p + w_0)}}{2w_4}$$

Given a probability s_p , the actual decision weight $d(s_p)$ is $\frac{1}{z(s_p)}$. Only a single solution is appropriate; in most cases the second solution is either negative or greater than 1. We set $d(1)$ to 1 and $d(0)$ to 0. If in any other case $d(s_p) > 1$ or $d(s_p) < 0$, we set it to 1 or 0 respectively (should not happen with sufficient data).

3.2 Assessing the value of multiple prospects

Using the decision weights obtained above, we can assess a value of a prospect. Given a prospect with a probability s_p of winning an outcome s_y , using the decision weight $d(s_p)$ we obtain that the value of the prospect for the user is simply:

$$u(s_p, s_y) = d(s_p) \cdot s_y \quad (1)$$

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. Note that there may be up to $n = 2^k$ possible outcomes and they must be sorted.

Once we obtain the probabilities $\vec{p} = p^1, p^2 \dots p^n$ ($\sum_{i=1}^n p^i = 1$) and outcomes $\vec{y} = y^1, y^2 \dots y^n$ ($y^1 < \dots < y^n$), using the Cumulative Prospect Theory [11] we assess the user value.

3.3 Advice Provision

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. This search can be invoked using any search method (such as genetic algorithm or interior-point).

Since human full rationality assumption is broadly used and is a very common assumption, we compare the performance of SAMP to the performance of a *rational agent* which assumes human full rationality (i.e. assumes people want to maximize their expected monetary value). Performance is measured in terms of human satisfaction. In all selection problems the rational agent advises the user to choose the uncertain outcome if $s_p \cdot s_y \geq s_x$, and the guaranteed outcome otherwise. The rational agent's advice does not depend on the number of selection problems k . A user following the rational agent's advice, will maximize his expected outcome.

4. EXPERIMENTAL DESIGN

We ran our experiments using Amazon's Mechanical Turk. We recruited 52 participants for SAMP's learning phase and 202 participants for evaluating 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. 52% of the subjects were males and 48% were females. Subjects' ages ranged from 18 to 75, with a mean of 32.2 and median of 29. All subjects were residents of the USA.

We set $k = 5$, i.e. the subjects had to make their choice regarding five prospect selection problems. This number was chosen to assure that the subjects would find it difficult to bracket on their own. The guaranteed outcome was drawn uniformly between 2 and 10 cents, and the probability to win the uncertain outcome (s_p) was drawn uniformly from $\{0.05, 0.1, 0.15, 0.2, \dots, 0.95\}$. Choosing the uncertain outcome yield up to 35% more on the expected utility value of the guaranteed outcome.

Once the subjects selected their preferences, we invoked a lottery on each of the prospect selection problems and paid them according to their preferences. I.e. if a subject chose the guaranteed outcome for a certain selection problem he was paid accordingly, and if he chose the uncertain outcome, the system randomly generated a number r in $[0, 1]$ and paid the subject s_y if $r \leq s_p$. In order to encourage narrow bracketing in the learning phase, the subjects were told to answer each selection problem as if it were standalone, and instead of having the system invoke all the selection problems, the system randomly selected a single selection problem and invoked only that one. The subjects were fully aware of this process.

In the evaluation phase, the system obviously invoked all five prospect selection problems and the subjects were also presented with the agent's advice. 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, where following the exact advice given by the agent, guarantees a win of at least 8 cents, but offers up to 89 cents. The interpretation of the pie chart was explained to the subjects, and their comments indicated that they clearly understood it.




Figure 1: A visualization of the agent's advice

Probability	Decision Weight	Probability	Decision Weight
0.10	0.07	0.25	0.20
0.40	0.32	0.60	0.48
0.75	0.64	0.90	0.85

Table 1: Decision weights for selected probabilities

After submitting their choices and told the final results, 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, where 1 is associated with "not at all" in the first two questions, and "very bad advice" in the third question, 3 is associated with "o.k.", and 5 is associated with "very happy" in the first two questions, and "very good advice" in the third question.

4.1 SAMP construction

Recall that in SAMP construction, after the data is collected a logistic regression classifier is built using this data. This classifier is built in order to find the decision weights, therefore the classifier is trained on all the data. However, in order to evaluate the performance of the classifier we also ran a 10 fold cross validation on the data. We found that the accuracy of the classifier was 67.2% (F-Measure 0.67). These results are satisfying, as we are considering a general human model and people differ in their risk attitudes. Furthermore, this classifier allows us to build a good enough model for people's decision weights. Using other classifiers yield similar results (SVM - 67%, neural networks - 65.6%), however, logistic regression enables the model extraction described in section 3.1.

After extracting the parameters we were able to calculate the decision weights for any given probability. Several decision weights learned by SAMP for some selected probabilities can be found in Table 1. For example the decision weight associated with 0.25 is 0.20. This indicates that people should be indifferent between a prospect that offers a 25% chance to win 10 cents and a guaranteed outcome of 2 cents. Note that the decision weights are lower than the actual probabilities in all cases; this indicates that, in our environment, people were risk averse regardless of the probability.

Using these decision weights SAMP builds a general human model as described in section 3.2. Since we used a small k , we had a small search space and therefore SAMP used an exhaustive search when searching for the best combination to be advised to the users. Figure 1 presents an example of a SAMP's advice, in a scenario where the guaranteed outcome among all selection problems sums




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

up to 21 cents. Following the advice seems preferable to choosing the guaranteed outcome, as the probability of gaining less than the guaranteed outcome by following the advice is less than 14% and the user is very likely to gain more than the guaranteed outcome, with a slight chance of gaining much more. The expected utility from following SAMP's advice is 25.7 cents, however, SAMP evaluated the human value of its advice in 23.4 cents

4.2 Results

Figure 2 presents the final user 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. Multi-variant ANOVA test on all three parameters (together) indicates that these results differ significantly ($p < 0.05$). In the first two questions results differ significantly ($p < 0.05$) also using the student's t-test between the two groups, however, results for the third question only came close to the significance level by the student's t-test (with $p = 0.07$).

We also compare the fraction of prospect selection problems where each user followed the advice for that selection problem. The advice given by SAMP for each of the selection problems was followed 76.5% of the time, while the advice given by the rational agent was followed only 70.8% of the time. Since more users followed the advice given by SAMP, on average, users who received SAMP's advice resulted with a non-statistically significant higher outcome (34.5) than those who received the rational advice (32.2). We attach minor importance to this achievement, as SAMP is destined to increase the users' satisfaction and not expected outcome.

To sum up the results, the subjects' satisfaction level in all parameters was greater when they received SAMP's advice than when they received the fully rational agent's advice. This indicates that the assumptions used by SAMP are more accurate than assuming full rationality. However, the rational agent also performed well, since its advice was followed in a large majority of selection problems, and reached an average score of 3.36 when asked how good it was (3 is associated with "o.k."). Therefore, the rational agent may be good enough for a cold start (when no training data is available).

5. CONCLUSIONS AND FUTURE WORK

In this paper we present an agent which advises people in an environment which includes multiple prospect selection problems, where in each selection problem the user must choose between a guaranteed outcome and an uncertain outcome. This agent collects data on humans in the desirable environment and, uses it to build a general human model with prospect theory concepts. Using this

model, the agent searches for a combination of selections which is most favorable for humans and recommends it to the user. We present the resulted combination using a pie chart which visualizes the probability of each possible outcome. SAMP significantly outperforms a fully rational agent in terms of user satisfaction.

Reaching such an achievement using a general human model is conspicuous since people differ from each other, and many times advice which might be good for one user may not be as good for another. In this work we have proven the concept of using decision weights, showed how to learn them from data and generalize their use to a combination of multiple prospect selection problems.

Research has shown that people act differently when the stakes are higher, and tend to be more risk-averse [5]. Nonetheless, many users commented that they would act differently if the stakes were higher. Since people tend to reject favorable prospects when the stakes are higher, narrow bracketing will cause them to reject many more favorable prospects than they should. We therefore expect that, when the stakes are higher, our agent will have a larger impact on the user satisfaction. However, in our settings raising the stakes would not allow us to actually pay the subjects and the prospect selection problems would have needed to be hypothetical

In current work, the experiments were based on a single trial. We intend to test a scenario where the user doesn't see all of the selection problems at once, but must make his selections sequentially.

6. ACKNOWLEDGMENTS

This work was supported in part by ERC 267523, MURI W911NF-08-1-0144, ARO W911NF0910206 and W911NF1110344.

7. REFERENCES

- [1] A. Azaria, Z. Rabinovich, S. Kraus, C. V. Goldman, and Y. Gal. Strategic advice provision in repeated human-agent interactions. In *AAAI*, 2012.
- [2] U. Gneezy and J. Potters. An experiment on risk taking and evaluation periods. *Quarterly Journal of Economics*, pages 631–645, 1997.
- [3] G. Harrison and E. Rutstr m. Expected utility theory and prospect theory: one wedding and a decent funeral. *Experimental Economics*, 12:133–158.
- [4] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems (TOIS)*, 22(1):5–53, 2004.
- [5] S. J. Kachelmeier and M. Shehata. Examining risk preferences under high monetary incentives: Experimental evidence from the people's republic of china. *American Economic Review*, 82(5):1120–41, December 1992.
- [6] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.
- [7] D. Read, G. Loewenstein, and M. Rabin. Choice bracketing. *Journal of Risk and Uncertainty*, 19(1):171–197, Dec. 1999.
- [8] J. Rieskamp. The probabilistic nature of preferential choice. *Journal of Experimental Psychology*, 34:1446–1465, 2008.
- [9] D. Sarne, A. Elmalech, B. J. Grosz, and M. Geva. Less is more: restructuring decisions to improve agent search. In *Proc. of AAMAS*, pages 431–438, 2011.
- [10] A. Tversky and D. Kahneman. The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458, Jan. 1981.
- [11] A. Tversky and D. Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, Oct. 1992.