Detecting and Adapting to Change in User Preference

Angelo Restificar
School of Elec. Eng. & Comp. Science
Oregon State University
Corvallis, OR 97331

Peter Haddawy
CSIM Program
Asian Institute of Technology
Pathumthani, Thailand 12120

John Miyamoto
Department of Psychology
University of Washington
Seattle, WA 98195

Abstract

The development of automated preference elicitation tools has seen increased interest among researchers in recent years due to a growing interest in such diverse problems as development of user-adaptive software and greater involvement of patients in medical decision making. These tools not only must facilitate the elicitation of reliable information without overly fatiguing the interviewee but must also take into account changes in preferences. In this paper, we introduce two complementary indicators for detecting change in preference which can be used depending on the granularity of observed information. The first indicator exploits conflicts between the current model and the observed preference by using intervals mapped to gamble questions as guides in observing changes in risk attitudes. The second indicator relies on answers to gamble questions, and uses Chebyshev's inequality to infer the user's risk attitude. The model adapts to the change in preference by relearning whenever an indicator exceeds some preset threshold. We implemented our utility model using knowledge-based artificial neural networks that encode assumptions about a decision maker's preferences. This allows us to learn a decision maker's utility function from a relatively small set of answers to gamble questions thereby minimizing elicitation cost. Results of our experiments on a simulated change of real patient preference data suggest significant gain in performance when the utility model adapts to change in preference.

Introduction

Various applications ranging from development of user-adaptive software to patient-centered medical decision care require the ability to acquire and model people's preferences. In developing automated tools for preference modeling not only must reliable information be elicited without overly fatiguing the decision maker but also the issue of change in preference must be taken into consideration. In dynamic environments, the change in preference can be induced by several factors. Previous shopping experiences for instance could affect subsequent preferences. An agent may revise preferences when constraints and resources change, *e.g.,* when outcomes with respect to an income level decrease or increase (Binswanger 1981). Preferences could even change by some effect of predisposition (Gibbs 1997). In medical decision-making, preference change, also known as response shift (Schwartz & Sprangers 2000), can be triggered by a change in health status which may result in a change in the respondent's internal standards of measurement or values. Moreover, preferences may need to be revised to handle novel circumstances (Doyle & Thomason 1999; Doyle 2004).

An adequate mechanism must, therefore, be in place to detect preference change as well as facilitate minimal cost of elicitation. We address the problem of detecting change by using indicators that determine differences in risk attitudes. Our model adapts to the change by relearning whenever risk attitude indicators exceed some preset thresholds. More specifically, we show that the use of intervals to approximate risk behavior coincides with Keeney and Raiffa's characterization of risk attitude. We introduce a specific measure called *disagreement ratio* to compare new risk attitude information against existing information in the preference model and propose the use of conflict sets to detect inconsistency. Moreover, when certainty equivalents of gambles are available, differences in the average of certainty equivalents taken at different timepoints can be used to detect a possible change in risk attitude information.

Following previous work (Haddawy *et al.* 2003; Restificar *et al.* 2002), we use knowledge-based artificial neural networks or KBANN (Towell & Shavlik 1995) that encode assumptions about a decision maker's preferences to allow us to learn a decision maker's utility function from a relatively small set of elicited answers to gamble questions, thereby minimizing elicitation cost. The theory refinement approach to preference elicitation pioneered by Haddawy *et al.* offers several advantages over the traditional approaches (Keeney & Raiffa 1993). In theory refinement approach, no assumption about the structure of preferences, *e.g.,* utility independence, is made. More importantly, elicitation can start from a possibly incomplete and approximate domain theory which can then be made more accurate by training on examples. More accurate prior domain knowledge leads to a reduction in learning time and elicitation cost. The network which can represent both linear and nonlinear functions encodes the decision maker's utility function, which can be easily extracted by simulating elicitation, treating the network as the decision maker. Curve fitting techniques such as in Keeney & Raiffa (1993), can be used to recover the utility function based on answers to gamble questions queried from the network.

We performed experiments simulating change of preference based on real patient data. The set of data is from Miyamoto & Eraker (1988) which contains answers to gamble questions of 27 subjects. The gamble questions were about years of survival and quality of life. Our results suggest significant gain in performance when the utility model adapts to preference change.

Preliminaries

In this paper, we will assume that the utility function can be nonmonotonic. However, in such a case the range of the attribute is divided into partitions such that the utility function is monotonic in each partition. This assumption offers an operational advantage. Partitions in which the utility is monotonic guarantee unique certainty equivalent for each lottery or gamble question whose outcomes lie in the said partitions. This method of treating nonmonotonic utility functions so that the range of the attribute is monotonic in each partition is mentioned by Keeney & Raiffa (1993). For convenience, we will be using 'lottery' and 'gamble' interchangeably to mean the same thing.

Definition 1 *Let D be a domain, U be a utility function over D, and let o_i and o_j be outcomes in a gamble G where o_i occurs with a probability p, o_j occurs with a probability $(1-p)$, and $o_i, o_j \in D$. Let us denote this gamble G by the shorthand $(o_i; p; o_j)$. If $p = 0.5$, then G is called an even-chance gamble. A certainty equivalent, Y_{ceq}, is an amount $\hat{o}$ such that the decision maker (DM) is indifferent between G and $\hat{o}$. Thus, $U(\hat{o}) = pU(o_i) + (1-p)U(o_j)$ or $\hat{o} = U^{-1}[pU(o_i) + (1-p)U(o_j)]$. A proportional match PM corresponding to $\hat{o}$ of an even-chance gamble is defined as $PM = \frac{\hat{o}-o_i}{o_j-o_i}$, $o_j > o_i$.*

The concept of proportional match allows us to compare certainty equivalents for different pairs of o_i and o_j. The idea of certainty equivalent is also used in the following theorem to determine the user's or decision-maker's (DM) risk attitude.

Theorem 1 (Keeney & Raiffa, 1993) *For increasing utility functions, a decision maker is risk averse (resp. risk-prone) if and only if his certainty equivalent for any nondegenerate lottery is less than (resp. greater than) the expected consequence of that lottery.*

In practice, however, we can only ask a limited number of lottery questions, lest the DM experience fatigue and thus potentially give unreliable answers. The idea is to assess the DM's utility function and be able to detect changes to it when it has occurred.

Change in Risk Attitude

We will first consider the issue of revised preferences where change is characterized by a change in risk attitude from either a risk averse, risk prone, or risk neutral type to any of the other two types. The occurrence of such change can be observed in many decision making scenarios. For instance, the agents' actions in a business negotiation could be affected by factors like time constraints, increasing cost of communication, or a change in the availability of resources. Assume a vendor of a perishable item in a market without refrigeration facility. At the start of the day the vendor will unlikely yield to a lower price but when the day is ending he will likely be less risk-seeking if someone engages him in a negotiation over the price of the item. In medical decision making, a patient's risk attitude toward medical treatments can change as a result of change to the relative importance of an attribute's contribution toward a health state's desirability (Schwartz & Sprangers 2000). For example, time spent with the family may become increasingly more important to a cancer patient as health state worsens.

The basic idea of the technique proposed in this section is to partition the domain of the attribute Y into subintervals and map these subintervals into lottery questions so that the answers to the lottery questions characterize the risk attitude of the DM on the subintervals. We then monitor changes to the subintervals each time information is available. Sometimes we will refer to these subintervals as intervals when no confusion will arise.

Definition 2 *Let D be the domain over which attribute Y is defined. Let $I = [c, d]$ and $I' = [c', d']$ be intervals defined over D. The intervals I and I' are said to be **distinct** from each other if either their intersection is empty or the intersection contains a single point common to both. Otherwise, they are called **non-distinct**.*

The idea is to approximate the risk attitude of the DM on a given interval. So for example, if we want to know the DM's risk attitude over the entire domain of the attribute Y, we should ask lottery questions whose outcomes are endpoints of the domain. The definition below will allow us to ask the appropriate questions on a given domain.

Definition 3 *Let $C = \{[c, d] \mid c, d \in D,\ c \leq d\}$ and let $G \in \mathbf{G}$, $\mathbf{G} = \{(c; p; d) \mid [c, d] \in C,\ 0 < p < 1\}$. Define $\zeta(G) = [c, d]$. Let $G_1, G_2 \in \mathbf{G}$ where $G_1 = (c_{1_i}; p_1; d_{1_j})$ and $G_2 = (c_{2_i}; p_2; d_{2_j})$. G_1 and G_2 are said to be distinguishable whenever at least one of the following conditions hold: (a) $p_1 \neq p_2$ (b) $c_{1_i} \neq c_{2_i}$ (c) $d_{1_j} \neq d_{2_j}$.*

Definition 3 relates specific intervals in the domain of the attribute to gamble questions. Note also that the same interval is linked to all gambles $G \in \mathbf{G}$ with the same pair of outcomes regardless of the probabilities p associated with them. For our purpose of determining the risk attitude type from either risk averse, risk prone or risk neutral the use of the same interval is correct since by Theorem 1 if the DM is risk averse (resp. risk prone) then the certainty equivalent must be less than (resp. greater than) the expected consequence for *any* nondegenerate gamble.

Definition 4 *Let D be the domain over which attribute Y is defined. Let $\mathbf{I}$ be the set of all closed intervals in D and let $C \subseteq \mathbf{I}$. The set of intervals C **covers** R if $\bigcup_{I \in C} [c, d] = R$. Define C as a cover of any subset of intervals whose endpoints are in R. A cover that contains only distinct intervals is said to be a **strict cover**.*

We want an approximation procedure P to output the same risk attitude that can be inferred from currently available answers to gamble questions.

Definition 5 *Given answers to a finite number of gamble questions $G = (c; p; d)$, where $c, d \in D$, a procedure P approximates the DM's risk attitude* ***over a domain*** *D iff P returns risk-averse (resp. risk-prone, risk-neutral) when the DM's certainty equivalent is less than (resp. greater than, equal) the expected consequence for any nondegenerate G. Given answers to a finite number of gamble questions $G = (c; p; d)$, where $I = [c, d]$, P approximates the DM's risk attitude* ***over an interval*** *I of D iff P returns risk-averse (resp. risk-prone, risk-neutral) when the DM's certainty equivalent is less than (resp. greater than, equal) the expected consequence for any nondegenerate G. Let C be a cover for some region R in D. If the procedure P returns the same result for all $[c, d] \in C$ then C is said to be a* ***consistent cover****.*

We now define the function ρ that maps an interval to one of the risk attitude types. The image of an interval I under ρ depends on the relation between the certainty equivalent of the gamble G mapped to I under ζ and the expected value of the gamble G.

Definition 6 *Let $\hat{c}$ be the certainty equivalent of the gamble $G = (c; p; d)$ and $\mathbf{E}[G] = pc + (1 - p)d$. Let $\rho : \mathbf{I} \rightarrow \{RA, RP, RN\}$. Define*

$$\rho(I) = \begin{cases} \textit{RA} & \textit{if } \hat{c} < \mathbf{E}[G] \\ \textit{RP} & \textit{if } \hat{c} > \mathbf{E}[G] \\ \textit{RN} & \textit{if } \hat{c} = \mathbf{E}[G] \end{cases}$$

Furthermore, $\rho(C) = \{\rho(I) \mid I \in C\}$. Also, $\rho(C) =$ RA[1] (resp. RP, RN) iff $\forall\, I \in C$, $\rho(I) =$ RA (resp. RP, RN).

Algorithm 1 presents a simple approximation procedure sufficient for detecting changes in risk attitude over a domain in which the utility function is monotonic. It takes as input a cover C and returns the approximated risk attitude of the DM on a given region in domain D. This approximation is particularly helpful when there is only a limited number of questions that can be asked from the DM, as is usually the case in practice. Let G denote a gamble question and QUERY(G) be a user query that poses the question G to the decision maker DM. The function QUERY(G) returns the DM's certainty equivalent of the corresponding gamble. Let I_i be the ith interval in C and let the symbol $\perp$ mean inconclusive. Let $\mathbf{G}$ be a set of m distinguishable gamble questions $\{G_{i_1} \ldots G_{i_m}\}$, $m \geq 1$, for the interval I_i. Also, let $\rho(I_{i_j})$ be the risk attitude obtained using the answer to QUERY(G_{i_j}).

For each interval $I \in C$ the answer to the corresponding gamble question G is checked. If the set of answers is consistent then the procedure returns either risk-averse (RA), risk-prone (RP), or risk-neutral (RN). Otherwise, no conclusion can be reached and the procedure returns inconclusive $(\perp)$. Lines 5-6 guarantee that any inconsistency will be detected. The intuition is to approximate the risk attitude of the DM at regions of the domain by using only currently available information. Note that we are not restricted to ask gamble questions that correspond only to the subintervals

[1]We slightly abuse the notation and write $\rho(C) =$ RA instead of $\rho(C) = \{$RA$\}$.

Algorithm 1: Approximation algorithm (RAAA)
Input: A cover C of $I_1, \ldots, I_i, \ldots, I_n$
Output: $\rho(C)$
(1) **foreach** $I_i \in C$
(2) **foreach** $G_{i_j} \in \mathbf{G}$
(3) QUERY(G_{i_j})
(4) $\rho(C) \leftarrow \rho(C) \cup \{\rho(I_{i_j})\}$
(5) **if** $(\rho(C) \neq \text{RA}) \wedge (\rho(C) \neq \text{RP}) \wedge (\rho(C) \neq \text{RN})$
(6) **return** $\perp$
(7) **else**
(8) **return** $\rho(C)$

of a strict cover. Gamble questions maybe used to elicit approximations to different subintervals that overlap. The only condition for our approximation procedure to output a specific risk attitude is that C is a consistent cover of the set of (sub)intervals. Of course, it is possible for a domain that is partitioned by at least two intervals to be risk-averse in one interval and risk-prone in another. In this case, our approximation procedure neither says that the DM is risk-averse nor risk-prone in that domain. This can be interpreted as a nonuniform answer over the domain. The procedure concludes that the DM is risk-averse (resp. risk-prone, risk-neutral) in that domain if and only if all intervals in the cover are.

Theorem 2 *Let C be a cover for some region R over domain D. Let $\mathbf{G}$ be a set of n distinguishable gamble questions $G_1, \ldots, G_n$, $n \geq 1$. Let the interval width $w(G_j)$ be the difference between the upper interval limit and lower interval limit of $\zeta(G_j)$, $1 \leq j \leq n$. As $n \rightarrow \infty$, the Risk Attitude Approximation Algorithm (RAAA) returns risk-averse (resp. risk-prone, risk-neutral) iff according to Theorem 1 (Keeney & Raiffa 1993) the utility function is risk-averse (resp. risk-prone, risk-neutral).*

Theorem 2 states that as the number of gamble questions becomes infinitely large the Risk Attitude Approximation Algorithm returns the same charaterization of risk attitude as Keeney and Raiffa's theorem. This allows us to characterize the risk attitude of the DM by using intervals that cover some region R in D. Algorithm 1 detects preference change even if the data available is in the form of indirect user feedbacks, *i.e.,* not in the form of the numeric certainty equivalent which comes from an interview. An observed action that does not conform to one that is expected is a form of indirect user feedback. Suppose that a DM is expected to reject a procedure for a certain illness and that the procedure has 50% success rate. Assume further that if the procedure is successful the DM can expect an additional 5 years of his life and that if the procedure is not successful the number of years he will survive will decrease by 2 years. If the DM is expected to reject such a procedure in favor of the status quo but instead agrees to it, such observed action can be interpreted as a result of a risk attitude that does not conform to the current utility model. Note that this is inferred from an observation other than an answer to a gamble question.

In order to measure the change in risk attitude, we will need to measure the difference between the newly obtained risk attitude information and that which is already in the utility model.

Definition 7 *Let $C \subseteq \mathbf{I}$, $\rho(C)$ be the risk attitude information on C, and $\rho(I')$ be the new risk attitude information on interval I', where $I, I' \in C$. Let I and I' be non-distinct intervals. The* ***conflict set*** *of I' is $S(I') = \{I \mid \rho(I) \neq \rho(I'),\ I \in C\}$. Now, let $I^* = I \cap I'$. For any $I \in S(I')$, the* ***conflict factor*** *of I is $\gamma(I, I') = \frac{d^* - c^*}{d' - c'}$.*

The conflict set is basically the set of intervals whose risk attitude information disagrees with I'. The disagreement ratio of an interval I' is simply the fraction of I' that is in conflict with the current model, *i.e.,* it gives a measure of the severity of conflict between the currently observed information and that which is already in the model.

Definition 8 *Let $\delta(I')$ denote the* ***disagreement ratio*** *of I' with respect to the current utility model. Compute $\delta(I')$ as follows:*
Step 1. Partition $S(I')$ into $S_1 \ldots S_n$ such that each S_i is convex.
Step 2. Let $I_i = I' \cap S_i$. The ***conflict factor*** *of S_i is $\gamma(S_i, I') = \frac{d_i - c_i}{d' - c'}$. Compute $\delta(I') = \sum_i \gamma(S_i, I') = \frac{\sum_i (d_i - c_i)}{d' - c'}$.*
The ***average disagreement ratio*** *of a set H of intervals I_j is $\overline{\delta}(H) = \frac{1}{n} \sum \delta(I_j)$ where $n = |H|$.*

The disagreement ratio provides an efficient method for detecting change, filtering new information, and updating the utility model. It can be used to detect changes in preference, *i.e.,* a change is noted if new information disagrees with what is currently known, at a level exceeding some threshold τ. In addition, it can be used to filter new information such that no new information can be added when the disagreement level exceeds τ. Higher weight of information reliability can be assigned to newly elicited information so that intervals that disagree with the new information can be removed from the model to maintain the disagreement level below the threshold τ. Moreover, since the risk attitude information associated with the intervals are approximations, the ratio $\delta(I')$, expresses the likelihood that the information in the current utility model disagrees with the new information. Note that $0 \leq \delta(I') \leq 1$. The value of $\delta(I')$ is 1 when the information in the current utility model totally disagrees with the new information. $\delta(I')$ is 0 when no information in the current utility model disagrees with the newly acquired information.

Example 1 *Let $C = \{[0.2, 0.4], [0.3, 0.5], [0.5, 0.6]\}$ where $\rho(C) = \{\rho([0.2, 0.4]) = RA,\ \rho([0.3, 0.5]) = RA,\ \rho([0.5, 0.6]) = RP\}$. Suppose that observations indicate that $\rho([0.35, 0.55])$ is not consistent with either RA or RN. So, $I' = [0.35, 0.55]$. The conflict set $S([0.35, 0.55]) = \{[0.2, 0.4],\ [0.3, 0.5]\}$. There is only one partition, $S_1 = S([0.35, 0.55])$. $\gamma(S_1, I') = \delta(I') = \frac{0.5-0.35}{0.55-0.35} = 0.75$.*

In Example 1, the information associated with $[0.2, 0.4]$ and $[0.3, 0.5]$ disagrees with the risk attitude associated with $[0.35, 0.55]$. Since the disagreement covers 75% of the region covered by $[0.35, 0.55]$, the disagreement ratio is 0.75.

Example 2 *Let $C = \{[0.2, 0.4],\ [0.25, 0.4],\ [0.3, 0.5],\ [0.5, 0.6]\}$ where $\rho(C) = \{\rho([0.2, 0.4]) = RA,\ \rho([0.25, 0.4]) = RA,\ \rho([0.3, 0.5]) = RP,\ \rho([0.5, 0.6]) = RA\}$. Suppose that observations indicate that $\rho([0.3, 0.6])$ is not consistent with RA. $S([0.3, 0.6]) = \{[0.2, 0.4],\ [0.25, 0.4],\ [0.5, 0.6]\}$ where there are two partitions $S_1 = \{[0.2, 0.4],\ [0.25, 0.4]\}$ and $S_2 = \{[0.5, 0.6]\}$. So $I_1 = [0.3, 0.4]$ and $I_2 = [0.5, 0.6]$. Hence $\delta([0.3, 0.6]) = \gamma(S_1, I') + \gamma(S_2, I') = \frac{(0.4-0.3)+(0.6-0.5)}{0.6-0.3} = 0.67$.*

Example 3 *Consider H and $\rho(H)$ that are the same as C and $\rho(C)$ in Example 2, respectively. $\delta([0.25, 0.4]) = \frac{0.4-0.3}{0.4-0.25} = 0.67$ since $S([0.25, 0.4]) = \{[0.3, 0.5]\}$. Similarly, $\delta(0.2, 0.4) = \frac{0.4-0.3}{0.4-0.2} = 0.50$. $\delta([0.5, 0.6]) = 0$ since no interval conflicts with $[0.5, 0.6]$. Since $S([0.3, 0.5]) = \{[0.25, 0.4], [0.2, 0.4]\}$ then $\delta([0.3, 0.5]) = \frac{0.4-0.3}{0.5-0.3} = 0.50$. Thus, $\overline{\delta}(H) = \frac{0.67+0.50+0+0.50}{4} = 0.42$.*

We now present an algorithm for detecting and updating risk attitude information given a user-defined threshold τ. τ specifies the level of disagreement that risk attitude information between intervals is allowed to have. Hence, a wide spectrum of 'noise' levels can be accommodated depending on the threshold. Let $S(I')$ be a conflict set and τ be a user-defined threshold, $0 \leq \tau \leq 1$. Let $\varphi : \mathbf{I} \rightarrow T$ be a timestamp function which marks the position of the interval I in a sequence of observations. All else being equal, the oldest information is removed from the model first.

Algorithm 2: Detecting risk attitude change
Input: C, τ, and an observation $\rho(I')$
Output: a possibly updated C
(1) Compute the conflict set $S(I')$ given $\rho(I')$
(2) Determine $\gamma(S_i, I')$ for each partition S_i of $S(I')$
(3) **while** $\delta(I') \geq \tau$
(4) Choose the partition S_i with the greatest $\gamma(S_i, I')$
(5) Remove the interval I with the greatest $\gamma(I, I')$ from S_i using φ to break any tie
(6) Update $\gamma(S_i, I')$

Algorithm 2 detects change in risk attitude using a user-defined threshold, τ. When a conflict exists between new and old data, it puts a higher weight of reliability to new observations and removes older information to keep the disagreement ratio within τ.

Example 4 *Consider C and $\rho(C)$ as in Example 2. Assume that $\tau = 0.5$. Since $\delta([0.3, 0.6])$ exceeds the threshold we detect a change in risk attitude. Note that $\gamma(S_1, [0.3, 0.6])$ and $\gamma(S_2, [0.3, 0.6])$ are both $\frac{1}{3}$ so we can choose either of the corresponding partitions S_1 and S_2. Suppose we choose $S_2 = \{[0.5, 0.6]\}$. We can then remove the interval $[0.5, 0.6]$ from C. The new conflict set is now $S([0.3, 0.6]) = \{[0.2, 0.4], [0.25, 0.4]\}$ which forms only one partition. So, $I_1 = [0.3, 0.4]$. The new disagreement ratio* *is* $\delta([0.3, 0.6]) = \frac{0.4-0.3}{0.6-0.3} = 0.33$. *Since* $\delta([0.3, 0.6])$ *is no longer as large as the threshold,* $\tau = 0.5$*, the procedure stops.*

It is possible that the new C no longer covers the same domain that the original C covered. This is a trade-off that has to be taken in order to bring down the level of disagreement below the user-defined threshold. The cover can be repaired by eliciting more information from the DM to fill the gaps caused by the removal of intervals using the update procedure.

In nonmonotonic utility functions, not only can risk attitude change but also the preference order over attributes. To handle the latter case, we can augment the technique given above by inferring from indirect user feedbacks or by asking additional queries. Each time QUERY(G_{i_j}) in Algorithm 1 is made to the DM we can also check whether the DM's preference order over the endpoints of the interval has changed. For example, it could be that previously on interval $[c, d]$ it is the case that $c \prec d$, but after the query it turns out that on interval $[c, d]$ it is the case that $c \succ d$. If preference order has changed then we simply need to determine from the DM's answers which partitions have consistent preference order, *i.e.*, partitions in which the function is monotonic. The algorithms above can then be applied to each of these partitions.

Chebyshev's inequality as a probabilistic bound

An indicator of a change in the risk behavior of the decision maker is the average, $\overline{Y}_{ceq}$, of all the certainty equivalents, Y_{ceq}. Using hypothesis testing we can test whether the difference between two sample averages is significantly different from zero. If this is the case, then a change has occured. $\overline{Y}_{ceq}$ can be computed by converting certainty equivalents of gambles with different outcomes to proportional matches (see Definition 1). Moreover, we can give a probabilistic bound on the risk attitude of the current sample data. Intuitively, if $\overline{Y}_{ceq} < \mathbf{E}[(o_i; p; o_j)]\ \forall p, 0 < p < 1$, then the DM is risk-averse. Of course, it could be the case that the DM is not risk-averse. However, based on the observed data the chances that he is risk averse would intuitively be higher than the chances of him being not. Without knowing the distribution of Y_{ceq}, we can use the sample $\overline{Y}_{ceq}$ and the sample standard deviation, $\sigma_{\overline{Y}}$ to give a probabilistic bound on the DM's risk behavior using Chebyshev's inequality. In particular, knowing $\overline{Y}_{ceq}$ and $\sigma_{\overline{Y}}$ we are interested to know what the chances are of any Y_{ceq} falling within some $k\sigma_{\overline{Y}}$ from $\overline{Y}_{ceq}$.

$$P(|Y_{ceq} - \overline{Y}_{ceq}| \leq k\sigma_{\overline{Y}}) \quad \geq \quad 1 - \frac{1}{k^2} \tag{1}$$

Using Chebyshev's inequality in Equation 1 with $k = 3$, the probability that Y_{ceq} is within $3\sigma_{\overline{Y}}$ is 0.89. If for instance,

$$\overline{Y}_{ceq} + 3\sigma_{\overline{Y}} \quad < \quad \mathbf{E}[(o_i; p; o_j)] \tag{2}$$

then intuitively, the Y_{ceq} we expect to get 89% of the time would be less than $\mathbf{E}[(o_i; p; o_j)]$. This implies that our observed data gives evidence that the DM is risk averse, based on both $\overline{Y}_{ceq}$ and $\sigma_{\overline{Y}}$.

The Chebyshev inequality gives us probabilistic bounds of Y_{ceq}. However, there are limitations in using this indicator alone. First, the measure depends on numerical values of Y_{ceq} which implies that it can not be used to detect preference change using indirect user feedback only. Second, the number of lottery questions whose consequences are within a specific subinterval could be very small. The measure would only be more reliable at regions where there is sufficient number of samples. Finally, the results of this indicator do not necessarily coincide with Keeney and Raiffa's theorem, *i.e.*, we maybe able to determine that there is a high probability that a region or the entire attribute domain is either risk-averse, risk-prone, or risk-neutral but this does not give us the assurance that the certainty equivalent of *any* nondegenerate lottery is less than, greater than, or equal to the expected consequence.

Data that comes from answers to gamble questions elicited within the same time frame maybe considered as part of a batch. In this case, change in risk attitude can be detected by comparing the $\overline{Y}_{ceq}$ for each batch. Moreover, a newer batch with a high disagreement ratio against a previous batch may also indicate change in risk attitude. In cases where data is in the form of continuous stream, a window adaptive heuristic similar to Widmer & Kubat (1996) can be used. A stable set of data can be viewed as a window of training instances whose $\sigma_{\overline{Y}}$ and average disagreement ratio do not exceed some specified thresholds.

Experiments and Results

We used the set of data from Miyamoto & Eraker (1988) which contains the answers of 27 subjects who were asked 24 even-chance gamble questions concerning years of survival and the quality of life. We evaluated our approach using the KBANN representation of a domain theory. The KBANN structure is as shown in Figure 1. (We urge the interested reader to see Haddawy *et al.* (2003) and Restificar *et al.* (2002) for more detailed discussion on the use of KBANN to model utilities.) The domain theory consists of rules that encode general assumptions about monotonicity and risk. In the rules that comprise the domain theory, let $Y_i, i = 1, 2$ be the duration of survival in number of years and Q be the health quality. Let Y_1 and Y_2 be the outcomes in an even-chance gamble G, $Y_1 > Y_2$, Y_{ceq} be the certainty equivalent of the gamble, and Y_{ce} be any arbitrary outcome for certain. Since Q is held fixed in each gamble question, it need not appear in the rules. The rules that comprise the domain theory are as follows:

$$(Y_1 > Y_{ce}) \wedge (Y_2 \geq Y_{ce}) \longrightarrow (G \succ Y_{ce}) \tag{3}$$

$$(Y_1 \geq Y_{ce}) \wedge (Y_2 > Y_{ce}) \longrightarrow (G \succ Y_{ce}) \tag{4}$$

$$Y_{ce} \leq ((Y_1 - Y_2)PM_{lo} + Y_2) \longrightarrow (G \succ Y_{ce}) \tag{5}$$

$$Y_{ce} > ((Y_1 - Y_2)PM_{hi} + Y_2) \longrightarrow (Y_{ce} \succ G) \tag{6}$$

Rules 3 and 4 say that whenever the outcomes of the even-chance gamble dominate Y_{ce}, the decision-maker prefers the gamble. Rules 5 and 6 encode attitude toward risk. Note that $[((Y_1 - Y_2)PM_{lo} + Y_2), ((Y_1 - Y_2)PM_{hi} + Y_2)]$ is another way of writing Y_{ceq} in interval form. This representation allows flexibility when approximating Y_{ceq}. Rule 5 says that

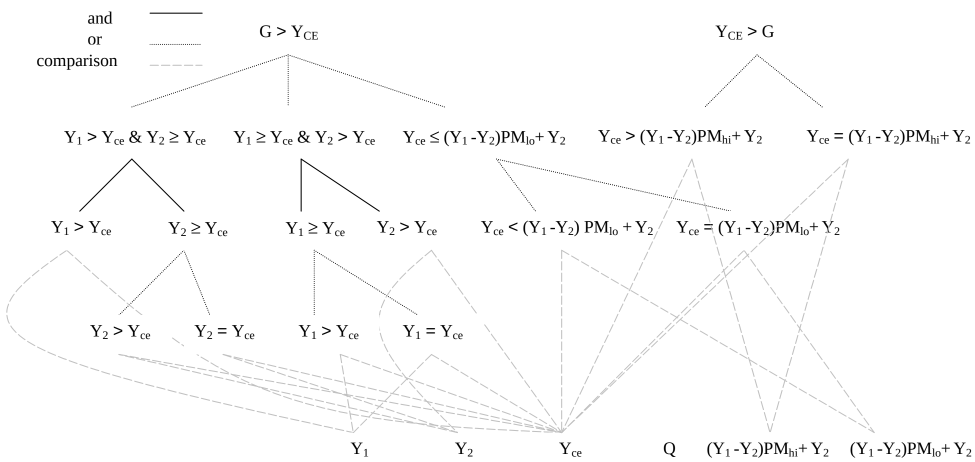

Figure 1: KBANN for the domain theory

a decision-maker prefers to gamble if $Y_{ce} \leq Y_{ceq}$. On the other hand, Rule 6 says that Y_{ce} is preferred whenever $Y_{ce} > Y_{ceq}$.

In order to perform experiments on preference change we used Miyamoto & Eraker's data from different individuals to simulate an individual's change of preference. We simulate a change from one risk attitude type to another by using $\overline{Y}_{ceq}$. Specifically, we chose subjects whose $\overline{Y}_{ceq}$ is either greater than 0.55 (P) or less than 0.45 (A), with $\sigma_{\overline{Y}} < 0.20$. Note that this does not necessarily mean that the subjects are risk-averse if $\overline{Y}_{ceq} < 0.45$ or risk-prone if $\overline{Y}_{ceq} > 0.55$. Using the probability bounds provided by the Chebyshev inequality the value of $\overline{Y}_{ceq}$ is just an indicator and since our representation does not assume any form of a family of equations, we can learn and approximate any kind of utility function. Although in our simulation we have assumed change from one risk attitude type to another on the entire domain over time, in cases where different attitude types might exist in different partitions of the same domain, as suggested for instance by Verhoef, Haan, & van Daal (1994), the techniques above can be applied to each respective partition.

We simulate an individual's preference change by using a sequence of six batches of data half of which has $\overline{Y}_{ceq} > 0.55$ and half of which has $\overline{Y}_{ceq} < 0.45$. In this context, the use of a batch is natural since answers provided by a single patient in the data are from a set of 24 gamble questions elicited within the same time period. In particular, we investigated the following sequences: A-A-A-P-P-P and P-P-P-A-A-A. These sequences simulate a change of preferences from a possibly risk-averse or risk-prone utility function to the other. The sequence A-A-A-P-P-P means that at time $t = 1$, KBANN learns using a data set whose $\overline{Y}_{ceq}$ is less than 0.45. At $t = 2$ and $t = 3$, KBANN takes as input two data sets where each $\overline{Y}_{ceq}$ is less than 0.45. At $t = 4, 5, 6$, KBANN takes as input a sequence of three data sets where $\overline{Y}_{ceq}$ is greater than 0.55. Similar meaning applies to the sequence P-P-P-A-A-A. The sequences A-A-A-P-P-P and P-P-P-A-A-A allow us to make observations regarding automatic detection and adaptation when drastic change occurs. The sequences also exhibit slight changes between succeeding data sets where relearning only happens when indicators exceed a set of thresholds. For this particular experiment, we adapt to the change in preference by relearning the network weights when at least one of the following conditions hold: (1) the average disagreement ratio for all new intervals exceed some threshold τ, (2) the difference between the average certainty equivalents of the current time frame and the last time frame the network has learned from scratch, exceeds some threshold β.

For convenience, we specifically chose answers from subjects that were asked the same set of standard gambles. For all 24 standard gamble questions, these subjects were asked 6 different questions, *i.e.*, each question was repeated 4 times. We therefore have 6 intervals from which to monitor change of preference. To smooth out the noise within the 6 intervals we took the average of the 4 answers in each interval and assumed it as representative of the DM's risk attitude at that interval. Data is available to the system in batches of 24 observations. This means that the system reads the answers to 24 standard gamble questions before it decides whether to use the current weights of the network or relearn from scratch. When relearning occurs the network only uses the data from current time frame. The weights of the network are those corresponding to the best network accuracy level found using 24-fold cross validation. In the succeeding time frames, if the system decides to adapt, the network weights are discarded, and the best set of weights found will replace the old set. Otherwise, the set of weights used to evaluate a succeeding time frame is the same as the weights used in the time frame before it. In order to compare the network performance between distinct thresholds we use as benchmark the performance of the network that learns at each time frame.

Figure 2 shows the performance of KBANN with and without detection at six time points of a simulated prefer-

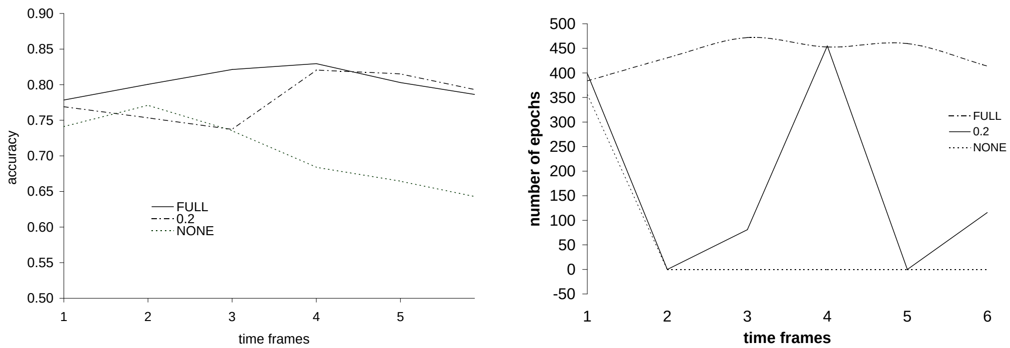

(a) Network accuracy

(b) Network training times

Figure 2: Adaptive and non-adaptive network performance and training times

ence change averaged over 10 runs. The top curve indicated by `FULL` shows the performance when there is relearning at each time frame for sequences `A-A-A-P-P-P` and `P-P-P-A-A-A`. The dotted bottom curve marked `NONE` shows the network accuracy when no adaptation is performed even when changes have been detected in time frames 2 to 6. The only data the network has learned from is that of the first time frame. The solid curve marked 0.2 describes the network accuracy in the case where the network relearns when the difference between the $\overline{Y}_{ceq}$ of the current time frame and that of the last time frame the network has learned from scratch exceeds by $\beta = 0.2$. The value of β can be anything between 0 and 1, inclusive. For our particular experiment, we have arbitrarily chosen the value of β where it is set at 0.2. For all data sets allowed into the model, the average disagreement ratio threshold τ is 0.99. This means that as long as no absolute inconsistency occurs, *i.e.*, it is not the case that the average disagreement ratio is 1, the risk attitude information from any interval in the data set is allowed into the model.

In Figure 2(a), there is a significant performance disparity beginning at the fourth time frame between the curve marked `FULL` and the curve marked `NONE`. The drastic change beginning at the fourth time frame caused the prediction accuracy to drop by at least 10% when the network does not adapt to the new risk attitude information. The curve marked 0.2 adapts to the information at the fourth time frame. Notably, the network accuracy has increased to about the same level as the `FULL` curve. From the results in Figure 2(a), the need for relearning becomes critical when the risk attitude evidence changes from either `A` to `P` or from `P` to `A`. Moreover, it is clear that there is a trade-off between training times and accuracy as indicated in Figure 2(b). Hence, a user-defined threshold such as τ and β can be used to choose a convenient tradeoff point.

Summary

In this paper, we presented two complementary methods for detecting and adapting to user preference change. The first method detects change between risk attitude types while the second method detects change within the same risk attitude type. One or both can be applied depending on the granularity of available information. We introduced the notion of disagreement ratio which provides a measure of the severity of conflict between the current model and the new information. It also provides a wide spectrum of control over the level of 'noise' that may be allowed into the model. The second indicator which uses Chebyshev's inequality is useful whenever the form of data received by the modeling component is in terms of answers to gamble questions. We have used knowledge-based artificial neural networks to represent and learn changing user preferences and results of our experiments suggest significant gain in performance whenever the utility model adapts to change in user preference.

References

Binswanger, H. P. 1981. Attitudes toward risk: Theoretical implications of an experiment in rural india. *The Economic Journal* 91(364):867–890.

Doyle, J., and Thomason, R. H. 1999. Background to qualitative decision theory. In *AI Magazine*. AAAI.

Doyle, J. 2004. Prospects for preferences. *Computational Intelligence* 20(2).

Gibbs, B. J. 1997. Predisposing the decision maker versus framing the decision: A consumer manipulation approach to dynamic preference. *Marketing Letters* 8(1):71–83.

Haddawy, P.; Ha, V.; Restificar, A.; Geisler, B.; and Miyamoto, J. 2003. Preference elicitation via theory refinement. *Journal of Machine Learning Research* 4:317–337.

Keeney, R., and Raiffa, H. 1993. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge University Press.

Miyamoto, J. M., and Eraker, S. A. 1988. A multiplicative model of the utility of survival duration and health quality. *Journal of Experimental Psychology: General* 117(1):3–20.

Restificar, A.; Haddawy, P.; Ha, V.; and Miyamoto, J. 2002.

Eliciting utilities by refining theories of monotonicity and risk. In *Working Notes of the AAAI-02 Workshop on Preferences in AI and CP: Symbolic Approaches*. American Association for Artificial Intelligence.

Schwartz, C. E., and Sprangers, M. A. G. 2000. Methodological approaches for assessing response shift in longitudinal health-related quality-of-life research. In Schwartz, C. E., and Sprangers, M. A. G., eds., *Adaptation to Changing Health: Response Shift to Quality of Life Research*. American Psychological Association.

Towell, G. G., and Shavlik, J. W. 1995. Knowledge-Based Artificial Neural Networks. *Artificial Intelligence* 70.

Verhoef, L. C. G.; Haan, A. F. J. D.; and van Daal, W. A. J. 1994. Risk attitude in gambles with years of life: Empirical support for prospect theory. *Medical Decision Making* 14(2):194–200.

Widmer, G., and Kubat, M. 1996. Learning in the presence of concept drift and hidden contexts. *Machine Learning* 23(2):69–101.