Making Rational Decisions using Adaptive Utility Elicitation

Urszula Chajewska
Computer Science Department
Stanford University
Stanford, CA 94305-9010
urszula@cs.stanford.edu

Daphne Koller
Computer Science Department
Stanford University
Stanford, CA 94305-9010
koller@cs.stanford.edu

Ronald Parr
Computer Science Department
Stanford University
Stanford, CA 94305-9010
parr@cs.stanford.edu

Abstract

Rational decision making requires full knowledge of the utility function of the person affected by the decisions. However, in many cases, the task of acquiring such knowledge is not feasible due to the size of the outcome space and the complexity of the utility elicitation process. Given that the amount of utility information we can acquire is limited, we need to make decisions with partial utility information and should carefully select which utility elicitation questions we ask. In this paper, we propose a new approach for this problem that utilizes a prior probability distribution over the person's utility function, perhaps learned from a population of similar people. The relevance of a utility elicitation question for the current decision problem can then be measured using its value of information. We propose an algorithm that interleaves the analysis of the decision problem and utility elicitation to allow these two tasks to inform each other. At every step, it asks the utility elicitation question giving us the highest value of information and computes the best strategy based on the information acquired so far, stopping when the expected utility loss resulting from our recommendation falls below a pre-specified threshold. We show how the various steps of this algorithm can be implemented efficiently.

Introduction

Rational decision making requires full knowledge of the utility function of the person affected by the decisions. According to traditional decision-theoretic principles, we should choose the sequence of decisions, or the *strategy*, which maximizes the expected utility (von Neumann & Morgenstern 1947). In order to compute the expected utility of a strategy, we need to know both the probabilities of all possible events in the decision problem and the utilities of all world states, or *outcomes*. The task of acquiring the probabilistic information is well understood. We know how to elicit such knowledge from experts or learn it from data.

Eliciting utility information is inherently harder. First, every person we advise may have a different utility function. Therefore, we need to elicit utility information not once, but many times, once for each user. Second, the task of utility elicitation is cognitively difficult and error prone. There are many elicitation techniques and the fact that they produce very different results when applied to the same person is well documented (Fromberg & Kane 1989). People find it hard to answer utility elicitation questions; they need to be trained beforehand and a significant percentage of them still give inconsistent answers. Finally, for many interesting, real-life decision problems, the outcome space is very large. There is a limit to the number of questions we can ask a person before fatigue will start playing a role. A decision tool requiring an interview several hours long is not likely to be widely used.

In order to apply decision-theoretic tools to such situations, we have to address two new issues. First, we need to find a way to make optimal or nearly-optimal decisions based on incomplete utility information. Second, in order to use the time and attention that our users are willing to give us well, we should also carefully choose the questions we ask to elicit utilities. In this paper we argue that the decision analysis and utility elicitation should not be considered to be two separate tasks, but rather two parts of one process. Each of these parts can influence and inform the other and together they make the decision making process more accurate and more efficient.

Our approach to utility elicitation and decision making departs from traditional approaches in one very important way: we treat utility as a random variable that is drawn from a known distribution. This assumption is often quite reasonable: many medical informatics centers collect large databases of utility functions for various decision problems for the purpose of cost-benefit analyses of new treatments. We can use such databases of utility functions to estimate the distribution of utility functions in the population (Chajewska & Koller 2000) and then use this estimate as a prior when we elicit utilities from the new users we advise.

This idea is key to our approach. First, it tells us how to choose a strategy relative to a partially specified utility function: we simply act optimally relative to our uncertain information about this random variable. Second, it provides a clear metric for evaluating different possible utility elicitation questions: the extent to which a question helps reach the optimal decision is simply its *value of information* (Howard 1966). This insight provides us with the basis for our algorithm. At every step, our algorithm computes the optimal strategy based on the current information. It then asks the elicitation question with the highest value of information, and the user's answer is incorporated into the model. This

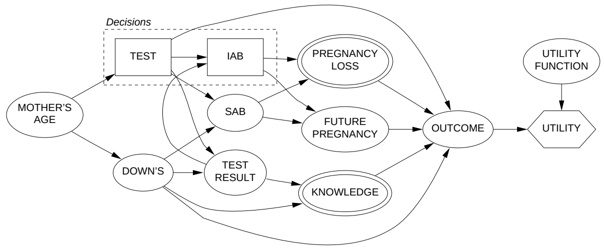

Figure 1: The decision model for prenatal diagnosis (courtesy of Joseph Norman, Stanford Medical Informatics).

process continues until the expected utility loss — the expectation, taken over the possible utility functions for the user, of the difference between the utility of using the current recommendation and the utility of using the strategy optimal for that user — falls below a pre-specified threshold. We provide a clean and efficient algorithm for implementing this scheme, despite the fact that it must deal with distributions over continuous utility variables.

Decision models

A sequential decision problem consists of several decisions, taken in sequence, typically with some information revealed between one decision and the next. Such a decision problem is often represented as an influence diagram (Howard & Matheson 1984). In our experiments, we will use a simplified version of a decision model developed at Stanford Medical Informatics for the task of prenatal testing. The influence diagram is shown in Figure 1. Prenatal testing is intended to diagnose the presence of a chromosomal abnormality such as Down's syndrome in the early weeks of pregnancy. The probability of such abnormality increases with maternal age. The two tests currently available to diagnose it, chorionic villus sampling (CVS) and amniocentesis (AMNIO), carry a significant risk of miscarriage above the baseline rate. The risk is higher for CVS, but it is more accurate and can be performed earlier in the pregnancy. In addition, both miscarriage (SAB) and elective termination of the pregnancy (IAB) may reduce the chances of future pregnancy. The influence diagram contains two decision nodes — which test to perform (CVS, AMNIO, or none), and whether to perform an elective termination. The edges going into the decision nodes represent information that is available to the decision maker before making the decision; for example, the test result (if any) is known before the decision about the elective termination.

An influence diagram can be solved by converting it into a *decision tree*, where at some nodes, the "choice" is made by nature and at others by the decision maker (Pearl 1988). We order the decisions $D_1, \ldots, D_k$, and partition the chance nodes into mutually exclusive and exhaustive sets $\mathbf{Y}_1, \ldots, \mathbf{Y}_k, \mathbf{Y}_{k+1}$, where $\mathbf{Y}_i$ are the variables that are revealed to the decision maker prior to D_i. The initial choice is made by nature about the random variables in $\mathbf{Y}_1$, with distribution $P(\mathbf{Y}_1)$. The next choice is made by the agent about D_1; the agent can make different choices about D_1 in different nodes in the tree, thereby allowing his decision to depend on the values of $\mathbf{Y}_1$. The tree continues in this way, until at the very end, nature decides on the remaining (unobserved) variables $\mathbf{Y}_{k+1}$, and the final utility is determined accordingly. The optimal strategy in the tree can be computed by the standard *expectimax* algorithm, which traverses the tree bottom up, with the agent doing a Max operation at each decision node, and nature doing an expectation operation at each chance node. While more sophisticated algorithms for influence diagram inference exist (Jensen, Jensen, & Dittmer 1994), they are essentially equivalent in their behavior to this simple decision-tree algorithm when applied to standard influence diagrams (those in *Howard normal form*).

Utilities

We begin with a review of some basic concepts in utility theory. Our presentation is based on concepts discussed in detail in (Keeney & Raiffa 1976).

Utility representation

Let O be the set of possible outcomes in our domain, $\{o_1, \ldots, o_m\}$. A utility function maps outcomes to real numbers. The naive representation of a utility function is a vector of real numbers, ascribing a utility to each possible outcome. Many real-life domains, however, involve fairly complex outcomes. In such cases, the space of outcomes is defined as the set of possible assignments of values to a set of relevant variables. The utility is a function of all of these values. This structure is fairly typical in complex decision problems. Thus, we define each outcome as an assignment to a set of attribute variables $\mathbf{X} = \{X_1, \ldots, X_n\}$. Each variable X_i has a domain $\mathrm{Dom}(X_i)$ of two or more elements. For example, in our domain, the outcomes include 'first trimester test (CVS), normal result, birth of a healthy baby', 'second trimester test (amniocentesis), procedure-related miscarriage, inability to conceive later', and 'no test, birth of a baby with Down's syndrome'. In general, the utility can depend on five attributes: T: type of testing; D: fetus status; L: possible loss of pregnancy; K: knowledge of the fetus's status; F: future successful pregnancy. The space of possible outcomes depending on these attributes is very large: $3 \times 2 \times 3 \times 3 \times 2 = 108$. Some of these can be eliminated as impossible, extremely unlikely, or absent in medical practice, but even then more than 70 remain.

In some cases, subsets of attribute variables exhibit some sort of independence, leading to structured utility representations. In particular, we might be able to represent the utility $U(T, D, L, K, F)$ as a sum of *subutility functions*, e.g., $U(T) + U(L, F) + U(K, L, D)$. In these cases, the utility function can be specified using a smaller number of parameters (27 in this example).

Utility elicitation

The fact that the utility function can be different for every patient involved, and the large number of parameters involved in specifying it, has often led clinical practitioners to use a single "universal" utility function for all patients. This approach has led, for example, to the uniform recommendation of prenatal testing for all women above the age of 35. This type of universal recommendation is rarely suitable for all women. In particular, our analysis of the model revealed the considerable influence of the utility function (especially the patient's attitude towards the risk of having a Down's child and toward a miscarriage) on the optimal choice of actions. An alternative approach is to elicit an individual utility function from each patient.

The *standard gamble* approach to utility elicitation (von Neumann & Morgenstern 1947) estimates the strength of a person's preference for an outcome by the risks he or she is willing to take to obtain it. Consider three outcomes o_1, o_2, and o_3 and a user with the preference ordering $o_1 \succ o_2 \succ o_3$. If he or she is offered a choice between o_2 for sure and a gamble in which o_1 will be received with probability p and o_3 with probability $(1-p)$, then, according to the theory, there exists a value of p for which the user will be indifferent. The outcome o_2 can then be assigned the utility value $pU(o_1) + (1-p)U(o_3)$.

The utility elicitation process starts by defining two universal best and worst outcomes, ($o_\top$ and $o_\perp$), which are then used as *anchors*, with values 1 and 0, to determine the relative utilities of all the other outcomes in the decision problem. A typical question in a medical domain is formulated in this way: "Imagine that you have a certain health condition which limits your activities in a specific way [the detailed description follows]. Imagine that there is a new experimental drug which only needs to be taken once. If taken, it will cure the condition p percent of the time, and $1-p$ percent of the time it will cause a painless death. Would you take the pill?" The question is repeated for a sequence of values of p until the user's preference changes. Several different ways of choosing the sequence can be used, including binary search, *ping-pong* (alternating between high and low values) and, most commonly, *titration* (reducing the value of p by a small amount). Recent research established that the final values are sensitive to the sequence choice (Lenert *et al.* 1998). Obviously, the user's answers are the most reliable far from the actual indifference point and most error-prone immediately around it. The need to use a sequence of questions for every outcome greatly increases the length of an interview and thus reduces the number of outcomes we can ask our user to consider.

Uncertainty over utilities

The discussion above, as well as practical experience, shows that it is virtually impossible to elicit a person's exact utility function. In this paper, we propose to circumvent this goal entirely. Rather than aiming at a completely specified utility function for a given patient, we maintain a *probability distribution* representing our beliefs about that patient's utility function. In other words, we view the different quantities $\{U_o\}_{o \in O}$ as a set of continuous-valued random variables (in the interval $[0, 1]$), with a joint probability density function (PDF) $p(\mathbf{U})$ ($\mathbf{U} = \{U_{o_1}, \ldots, U_{o_m}\}$), representing our beliefs about the patient's utilities.

This type of PDF can be represented in many ways; our approach applies to any representation that allows random samples to be generated from the PDF, thereby allowing moments of the distributions and expectations over it to be estimated numerically. However, our algorithm can be made more efficient in cases where the PDF allows some computations to be done in closed form, in particular PDFs that are jointly Gaussians or a mixture of Gaussians (cut off to fit in the $[0, 1]$ hypercube). A Gaussian can represent dependencies between a person's utilities for different outcomes. A mixture of Gaussians can represent distinct clusters in the population, whose utility functions are very different; in general, any PDF can be approximated arbitrarily well with a mixture of Gaussians (Bishop 1995). Furthermore, there are efficient algorithms for estimating these PDFs from data.

In cases where the utility function can be assumed to have some structure, as described above, it is better to represent and learn the distribution over utilities via a PDF over the (much smaller set of) parameters of the subutility functions (Chajewska & Koller 2000). As the utility variables are linear in these parameters, a mixture of Gaussians over the subutility parameters induces a mixture of Gaussians over $\mathbf{U}$.

The algorithm

Our approach is based on an integrated algorithm for decision making and utility elicitation. The answers to our utility elicitation questions inform the decision making procedure, the results of which help us select the most informative next utility elicitation question to ask.

When the system encounters a new patient, the only information available about her utility function is the prior PDF $p(\mathbf{U})$. The algorithm then cycles through the following steps.

- It computes the optimal strategy π relative to the current PDF $p(\mathbf{U})$.
- If this optimal strategy meets the stopping criterion, it stops and outputs π.
- Otherwise, it selects a utility elicitation question to ask the user, and asks it.
- It conditions $p(\mathbf{U})$ on the response.

Decisions under utility uncertainty

The first question we must address is how we make optimal decisions given uncertainty over the user's utility function $\mathbf{U}$. Fortunately, the answer to this question is easy (and well-known). Consider a given strategy π and a PDF $p(\mathbf{U})$. The expected utility of π for a fixed utility function $\mathbf{u}$ is: $\mathrm{EU}_\pi(\mathbf{u}) = \sum_{o \in O} P(o \mid \pi) u_o$. The expected utility under $p(\mathbf{U})$ can easily be shown to be $\mathrm{EU}_\pi(p) = \sum_o P(o \mid \pi) \mathcal{E}_p[u_o]$. Hence, we can find the best strategy given $p(\mathbf{U})$ by running standard influence diagram inference using, as the utility value for outcome o, the mean of U_o under p.

In general, we can compute the mean of U_o under p by Monte Carlo sampling. However, under the assumption that p is a mixture of Gaussians, we can compute it much more efficiently by combining the closed form integral of $u \exp(-u^2)$ with the integral tables readily available for $\exp(-u^2)$. (We omit details.)

Utility elicitation questions

A distribution $p(\mathbf{u})$ can be updated with additional information about the user's utility function, elicited from the user. We consider questions that follow the standard gamble pattern: "Given the choice between outcome o for sure and a lottery which gives $o_\top$ with probability s and $o_\bot$ with probability $1-s$, which will you choose?" We translate the response to this question to a constraint of the form $U_o < s$ or $U_o > s$, depending on the response. We call the value of s a *split point*. Consider a cycle of this process. Initially, we have a PDF p over the user's utilities. Let μ be the mean of $\mathbf{U}$ under p, and π^* the strategy that is optimal relative to μ. Now, consider a question regarding an outcome o and a split point s. If the user responds that $U_o < s$, we condition our PDF p, resulting in a new PDF $p_{<s}$; this will give us a new mean $\mu_{<s}$, and as a result, a new optimal strategy $\pi^*_{<s}$. Similarly, if she responds that $U_o > s$, we obtain a PDF $p_{>s}$ with $\mu_{>s}$ and associated optimal strategy $\pi^*_{>s}$. As further questions are asked and more information is obtained, our probability distribution p is updated, and the choice of optimal strategy changes to better fit the user's true preferences.

Note that our questioning pattern differs from standard gamble in a significant way: we do not ask about the same outcome for different values of p, until the indifference point is reached. Rather, we choose questions so as to reduce the total number of questions we need to ask the user. A given question will often be for a different outcome than the previous one.

We discuss the process of selecting the best question below.

Stopping criterion

After a sequence of utility elicitation questions, we will have a posterior distribution $\tilde{p}(\mathbf{U})$ over the user's utility function, and an associated candidate optimal policy $\hat{\pi}$. We would like to estimate the regret associated with stopping the utility elicitation and recommending $\hat{\pi}$. Assume that the user's true utility function is $\mathbf{u}$, and that the associated optimal strategy is $\pi^*_{\mathbf{u}}$. Then the user's loss is the difference between her expected utility, under $\mathbf{u}$, of $\pi^*_{\mathbf{u}}$, and her expected utility under the recommended strategy $\hat{\pi}$. The expected loss is the expectation of the loss under $\tilde{p}(\mathbf{u})$:

$$\int [\mathrm{EU}_{\pi^*_{\mathbf{u}}}(\mathbf{u}) - \mathrm{EU}_{\hat{\pi}}(\mathbf{u})]\tilde{p}(\mathbf{u})d\mathbf{u}.$$

While computing this integral exactly is impractical — we would need to compute the regions in which every strategy is optimal — we can approximate it quite easily using Monte Carlo methods. We simply sample utility functions $\mathbf{u}$ from $\tilde{p}$, use the influence diagram to compute the optimal strategy $\pi^*_{\mathbf{u}}$, and compute the loss for $\mathbf{u}$.

We can bound the number of samples needed using the upper bound on the worst case loss x, the desired threshold for expected utility loss ϵ and the confidence parameter δ by using Chebyshev's inequality: $N > \frac{x^2}{2\epsilon^2\delta}$.

Choosing the next question

One of the important advantages of explicitly modeling our uncertainty over the user's utility is that we obtain a simple metric for evaluating possible questions — the *value of information* measures the expected improvement in our decision quality derived from the answer to a question "is $U_o > s$?" We define the *posterior expected utility* after asking this question as:

$$\mathrm{PEU}(o,s) = \mathrm{EU}_{\pi^*_{<s}}(\mu_{<s})P(U_o < s) + \mathrm{EU}_{\pi^*_{>s}}(\mu_{>s})P(U_o > s)$$

This is an average of the expected utilities arising from the two possible answers to the question, weighted by how likely these two answers are. The value of information is this expression minus the current expected utility. Our goal is to find the outcome U_o, and the splitting point s for that outcome, that achieves the highest value of information. We note that this metric only evaluates the *myopic* value of information for the question. The full value of information, which takes into consideration all possible future combinations of questions and answers, is, as usual, intractable.

We will start our analysis for the case in which the utilities of different outcomes are probabilistically independent, i.e., the different variables U_o are marginally independent in p. We then relax this assumption in a later section.

Discretizing the problem

The first problem we encounter is that the utility variables range over a continuous space, so that there are infinitely many potential split points for each outcome. Fortunately, it turns out that we can restrict attention only to a finite number of them.

Let π be some strategy, and consider $\mathrm{EU}_\pi(\mathbf{U})$ as a function of a single utility variable $U_{o'}$:

$$\begin{aligned}
\mathrm{EU}_\pi(u_{o'}) &= p(u_{o'})u_{o'} + \sum_{o\neq o'} \int P(o|\pi)p(u_o|u_{o'})u_o du_o \quad (1)\\
&= p(u_{o'})u_{o'} + \sum_{o\neq o'} \int P(o|\pi)p(u_o)u_o du_o\\
&= p(u_{o'})u_{o'} + \sum_{o\neq o'} P(o|\pi)\mathcal{E}_p[U_o],
\end{aligned}$$

where the second equality is due to our independence assumption about utility variables.

Hence, the expected utility of a given strategy π is a linear function of U_o. The value for the optimal strategy for this problem is, for each value of U_o, the maximum over all strategies π. Thus, it is a piecewise-linear, convex function of U_o. We say that a strategy is *viable for o* if it is optimal for some value of U_o. We say that a particular value s is an *intersection point* if there are two viable strategies π_1 and π_2 that achieve the same expected utility at s, i.e., $\mathrm{EU}_{\pi_1}(s) = \mathrm{EU}_{\pi_2}(s)$.

Proposition 1: *The split point with the highest value of information will occur at one of the intersection points.*

Proof: Consider a potential split point s, and let π^*_L be the optimal strategy for the distribution $p_{<s}$ and π^*_R be the optimal strategy for the distribution $p_{>s}$. Let s^* be the strategy intersection point where $\mathrm{EU}(\pi^*_L) = \mathrm{EU}(\pi^*_R)$. Let's further assume, without loss of generality, that $s^* < s$. We want to show that $\mathrm{PEU}(o,s) \leq \mathrm{PEU}(o,s^*)$. It is easy to verify, for any $a < b < c$ and any strategy π, that

$$\begin{aligned}&\mathrm{EU}_\pi(\mu_{[a,c]})P([a,c]) \\ &\quad= \quad \mathrm{EU}_\pi(\mu_{[a,b]})P([a,b]) + \mathrm{EU}_\pi(\mu_{[b,c]})P([b,c])\end{aligned}$$

We have that

$$\begin{aligned}
&\mathrm{EU}_{\pi_L^*}(\mu_{<s})P(U_o < s) + \mathrm{EU}_{\pi_R^*}(\mu_{>s})P(U_o > s) \\
&\quad= \quad \mathrm{EU}_{\pi_L^*}(\mu_{<s^*})P(U_o < s^*) \\
&\qquad\qquad + \mathrm{EU}_{\pi_L^*}(\mu_{[s^*,s]})P(U_o \in [s^*,s]) \\
&\qquad\qquad + \mathrm{EU}_{\pi_R^*}(\mu_{>s})P(U_o > s) \\
&\quad\leq \quad \mathrm{EU}_{\pi_L^*}(\mu_{<s^*})P(U_o < s^*) \\
&\qquad\qquad + \mathrm{EU}_{\pi_R^*}(\mu_{[s^*,s]})P(U_o \in [s^*,s]) \\
&\qquad\qquad + \mathrm{EU}_{\pi_R^*}(\mu_{>s})P(U_o > s) \\
&\quad= \quad \mathrm{EU}_{\pi_L^*}(\mu_{<s^*})P(U_o < s^*) \\
&\qquad\qquad + \mathrm{EU}_{\pi_R^*}(\mu_{>s^*})P(U_o > s^*)
\end{aligned}$$

where the inequality is due to the fact that π_R^* dominates π_L^* for every $u_o > s^*$, and therefore also for $\mu_{[s^*,s]}$. Now, consider the two strategies $\pi^*_{<s^*}$ and $\pi^*_{>s^*}$ that are optimal for the distributions $p_{<s^*}$ and $p_{>s^*}$ respectively. These are not necessarily π_L^* and π_R^*, because the mean of $p_{<s^*}$, say, might not fall in the part of the region where π_L^* is optimal. However, it is easy to show that $\pi^*_{<s^*}$ and $\pi^*_{>s^*}$ can only improve the posterior expected utility. The result follows. ∎

We only need to consider those strategy intersection points where the viable strategies $\pi^*_{<s^*}$ and $\pi^*_{>s^*}$ intersect at s^*; otherwise, as the proof shows, the strategy intersection point of these two strategies would have higher value of information. How many points like this are there? Suppose we have N optimal strategies. Let's imagine moving the potential split point s from left to right over the range of U_o. We can mark an interval boundary whenever the optimal strategy for the area to the left or the optimal strategy for the area to the right of our split point changes. Note that once a strategy on the left side changes, it cannot change back: the mean $\mu_{<s}$ of U_o increases monotonically as we widen the region on the left, and the expected utility for any strategy is a linear function of this mean. Hence, the linear functions for any pair of strategies can cross at most once. Similarly, once a strategy on the right side changes, it cannot change back. As each strategy is optimal on each side at most once, we have at most $2N$ intervals, and at most $2N-1$ candidate split points. Each of these is only feasible, of course, if it is also a strategy intersection point of the two corresponding strategies. Thus, we need to consider only $2N-1$ split points, rather than N^2. We can execute this process efficiently using a simple binary search procedure, which utilizes the fact that we can find intersection points analytically. We omit details for lack of space.

The number of optimal strategies

The result above suggests that the number of VOI computations required is linear in the number of viable strategies. At first glance, this result might not be very reassuring. After all, there is an enormous number of strategies: exponential in the size of the decision tree. Any computation which requires us to consider all of them is much too expensive in all but the most trivial of decision problems. Fortunately, the number of *viable* strategies is exponentially smaller than the total number of strategies. Indeed, we show that it is linear in the size of the decision tree corresponding to our influence diagram. Given that we need to traverse the decision tree every time we use the decision model for finding an optimal strategy, this cost is very reasonable.

Proposition 2: *The number of strategies that are viable for o is at most the number of nodes in the decision tree corresponding to the influence diagram.*

Proof: We prove this result by induction on the depth of the tree. For the base case, a tree of depth 0 consists of a single leaf, where we have only a single strategy. In this case, the number of nodes is 1, and the number of viable strategies is also 1. For the inductive case, consider a tree of depth $d+1$. Let k be the number of children of the root, and let ℓ_i be the number of nodes in the subtree corresponding to the ith child. By the inductive hypothesis, the number of viable strategies for the ith child is at most ℓ_i. There are now two cases. Either the root is a max node or an expectation node. In the first case, the expected utility function $\mathrm{EU}(U_o)$ is the maximum of the functions of the children. In the second case, it is a weighted average of the functions of the children, where the weights are the probabilities annotating the edges going out of the root. In both cases, it is easy to verify that the function at the root is a piece-wise linear function, with a number of segments which is at most the total number of segments in the combined functions. (Intuitively, the reason is that the combined function can change from one linear function to another only at a point where one of the constituent functions changes from one linear function to another.) ∎

Correlated outcomes

Until now in this section, we have assumed that the different utility variables U_o are independent in $p(\mathbf{U})$. Unfortunately, this assumption is too strong in many cases. Indeed, it is quite likely that a woman's utility for one outcome involving a Down's baby will be correlated with her utility for another outcome involving the same event. In this section, we consider the more general case of an arbitrary distribution p.

We begin by assuming that our prior $p(\mathbf{U})$ is a multivariate Gaussian with an arbitrary covariance matrix, constrained to lie within the $[0,1]$ hypercube. Clearly, our utility function distribution cannot be truly Gaussian since utility functions are constrained to lie within the normalized range. Nevertheless, a Gaussian can be a reasonable approximation since the probability mass that lies outside of the normalized utility range will generally be negligibly small.

We use a convenient property of multivariate Gaussians to apply the algorithm of the previous section with almost no modifications: Given any variable $U_{o'}$, the conditional means of the remaining variables are linear functions of $U_{o'}$. In other words, in Eq. (1), although we no longer have that

$p(u_o \mid u_{o'}) = p(u_o)$, we do have that $\int p(u_o \mid u_{o'}) u_o du_o = g(u_{o'})$ for some linear function g. Thus, when we are enumerating the viable strategies for outcome o', as described in the section "Choosing the next question", we replace the means of the other U_o variables with their (linear) conditional means. The resulting function $\mathrm{EU}_\pi(U_{o'})$ is still a linear function of $U_{o'}$, so that the rest of the analysis remains unchanged.

Unfortunately, this analysis is insufficient for our purposes. Most obviously, it does not apply to the case of a mixture of Gaussians, where the mean of one variable is no longer a linear function of the other. There is a more subtle problem, however: even if our prior distribution is a multivariate Gaussian, once we condition our distribution on some information $U_{o''} > s$, the resulting posterior is no longer a multivariate Gaussian, but rather a "strip" of one. It can be verified that the conditional mean for this distribution is not a linear function.

We address both these difficulties using a simple approximation. Let p be our current distribution, conditioned on all of the relevant information. Rather than computing value of information relative to p, we approximate p using a distribution $\hat{p}$ which is a multivariate Gaussian. We note that our stopping criterion is always computed relative to the correct conditional distribution p.

In the case where p is a multivariate Gaussian, we can use a simple trick to maintain our distribution $\hat{p}$ as additional information is obtained. The key idea is that any multivariate Gaussian $p(X_1, X_2, \ldots, X_n)$ can be decomposed as a univariate Gaussian over X_1 and a *linear Gaussian (LG)* $p(X_2, \ldots, X_n \mid X_1)$ which defines a multivariate Gaussian over $X_2, \ldots, X_n$ with mean $\boldsymbol{\mu}(x_1)$ which is a vector linear function of x_1 and a fixed covariance matrix (Shachter & Kenley 1989). If we condition X_1 on some evidence, and approximate the result on a Gaussian, the parameterization of the LG $p(X_2, \ldots, X_n \mid X_1)$ does not change. Hence, the best approximation to the joint PDF as a multivariate Gaussian, given evidence on X_1, can be found by finding the best approximation to $p(X_1)$ as a univariate Gaussian and leaving the LG unchanged. We can then regenerate a new approximate multivariate Gaussian $\hat{p}'$ by multiplying $p(X_1)$ and the LG. In this case, the variable $U_{o'}$ on which we have evidence plays the role of X_1.

We can extend this approach to the case of a mixture of Gaussians. We use the same idea of approximating the distribution as a multivariate Gaussian, and then using our algorithm above for finding the optimal split point relative to the approximation. We then use the information about the X_1 (the utility $U_{o'}$) to define a posterior distribution. The only difference is that, in this case, we must also use our information about X_1 to update the mixture weights. This can be done using a standard application of Bayes rule.

Experimental results

Our database consisted of 51 utility functions elicited from pregnant women considering prenatal diagnosis, collected for an earlier study (Kuppermann *et al.* 1997). We used five-fold cross-validation for experiments, estimating the distribution using four sets and testing on the fifth. We ran these tests separately for every possible value of mother's age. Due to the small size of the database, we assumed that the utility function does not change with age and used all functions in the database to run tests for all ages.

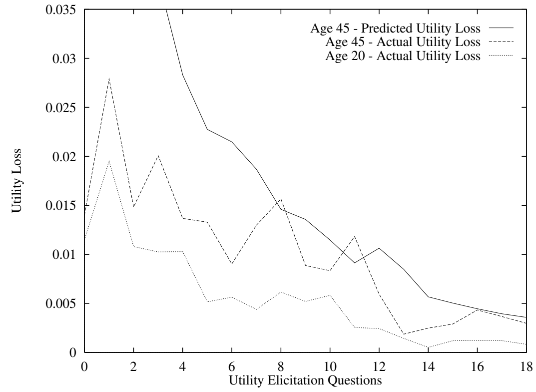

Figure 2: Expected and actual utility loss as a function of the number of qustions asked.

We present the results for an uncorrelated Gaussian; our current database is too small to allow reliable learning of more complex densities. Figure 2 shows the evolution of predicted and actual utility loss as more questions are asked. We see that the predicted utility loss starts out quite large, as the distribution is very broad. It gradually converges to the correct utility loss, and both gradually converge to zero. The predicted utility loss is usually an overestimate to the actual utility loss, implying that our algorithm is "safe" in not stopping prematurely. The overall results are summarized below; the ranges indicate behavior for different ages:

Number of questions asked

	avg	best	worst	std. dev.
$\epsilon = 0.05$	2.3–3.9	2	16	13–14.7
$\epsilon = 0.02$	5.9–9.0	3	16	12.1–30

Utility loss after last question

	avg	best	worst	std. dev.
$\epsilon = 0.05$	.01–.04	0	0.28	0.24–0.3
$\epsilon = 0.02$	.001–.015	0	0.087	0.05–0.14

Distance from indifference point

	Q1	Q2	Q3	avg (20 questions)
	0.41	0.21	0.13	0.21

As we can see, the number of questions asked is surprisingly small given the fact that we have 108 outcomes in the model. It increases slightly as we lower the threshold, but stays well within the bounds of what is possible in clinical practice. By comparison, the approach of (Chajewska *et al.* 1998) achieves an average of 7.55 questions for the same domain and an average utility loss of 0.035; with a much smaller number of questions, we achieve a utility loss which is comparable for $\epsilon = 0.05$ and substantially lower for $\epsilon = 0.02$. Furthermore, their approach provides no guarantees about the utility loss of the final recommendation. Finally, note that the split points our algorithm chooses are usually quite far from the indifference point, making the questions cognitively easy.

Related work

With recent advances in the power and scope of decision-theoretic systems, utility elicitation has become a lively and expanding area of research. A few projects are particularly relevant to this work. Jameson, et al. (1995) and Linden, et al. (1997) investigated the problem of eliciting partial utility models and reasoning with such models in case of certainty for domains such as shopping and making airline reservations on the Web. Poh and Horvitz (1993) discussed the value of the refinement of utility information. Finally, Jimison et al. (1992) suggested explicitly representing uncertainty about key utility parameters in the context of explaining medical decision models to the users.

Conclusion and extensions

We have presented a new approach for making decisions based on limited utility information, and for targeting our utility elicitation process so as to lead to a good decision using a small number of questions. Our approach is based on maintaining a probability distribution over possible utility values of the user, and using value of information for deciding which utility elicitation question will best help us in making rational decisions. We have presented algorithms for doing these computations efficiently.

We have presented results that suggest that our approach can make utility elicitation substantially easier for users of the decision model. Our results suggest that, in most cases, our chosen split point will not be in the immediate vicinity of the user's indifference point, thus making the task much easier cognitively. Furthermore, we have seen that our method substantially reduces the overall number of questions we have to ask before a good decision can be made; often, the number is as small as 2–3, with a very small utility loss. Indeed, one might expect that the overall decision quality will be better, because our method allows us to avoid errors resulting from the fatigue caused by the utility elicitation process.

There are many ways in which our work can be extended. For example, some questions are cognitively more difficult than others: questions near the indifference point are hard, a second consecutive question about the same outcome is cheaper than a question about an outcome discussed a few questions back, etc. It is easy to incorporate the cognitive cost of questions into the value of information computation.

A more general research direction is to exploit other aspects of our approach to representing uncertainty over utility functions. One possible application is to represent the dependence of this distribution on environmental variables. For example, it has been noted by the practitioners in the field that a woman who personally knows a Down's syndrome child is more likely to rate an outcome involving such a child as less desirable. A less obvious and more philosophically controversial application is to represent the probability that a user's preferences will change over time. We believe that the tools provided by probabilistic models — value of information, statistical learning, and more — will turn out to be extremely useful in this new type of modeling.

Acknowledgments We would like to thank Miriam Kuppermann for allowing us to use her data and Joseph Norman his model for the prenatal diagnosis domain. We are also grateful to Uri Lerner for his help in building the inference code for Gaussians and to Brian Milch for his influence diagram code. This research was supported by ARO under the MURI program "Integrated Approach to Intelligent Systems", grant number DAAH04-96-1-0341, by ONR contract N66001-97-C-8554 under DARPA's HPKB program, and through the generosity of the Powell Foundation and the Sloan Foundation.

References

Bishop, C. M. 1995. *Neural Networks for Pattern Recognition*. New York, NY: Oxford University Press.

Chajewska, U., and Koller, D. 2000. Utilities as random variables: Density estimation and structure discovery. In *Proc. UAI–00*. To appear.

Chajewska, U., Getoor, L., Norman, J., and Shahar, Y. 1998. Utility elicitation as a classification problem. In *Proc. UAI–98*, 79–88.

Fromberg, D. G., and Kane, R. L. 1989. Methodology for measuring health-state preferences—II: Scaling methods. *Journal of Clinical Epidemiology* 42(5):459–471.

Howard, R. A., and Matheson, J. E. 1984. Influence diagrams. In *The Principles and Applications of Decision Analysis*. Strategic Decisions Group.

Howard, R. A. 1966. Information value theory. *IEEE Transactions on Systems Science and Cybernetics* SSC-2:22–26.

Jameson, A., Schäfer, R., Simons, J., and Weis, T. 1995. Adaptive provision of evaluation-oriented information: Tasks and techniques. In *Proc. IJCAI–95*, 1886–1893.

Jensen, F., Jensen, F. V., and Dittmer, S. L. 1994. From influence diagrams to junction trees. In *Proc. UAI–94*.

Jimison, H. B., Fagan, L. M., Shachter, R. D., and Shortliffe, E. H. 1992. Patient-specific explanation in models of chronic disease. *AI in Medicine* 4:191–205.

Keeney, R. L., and Raiffa, H. 1976. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley & Sons, Inc.

Kuppermann, M., Shiboski, S., Feeny, D., Elkin, E. P., and Washington, A. E. 1997. Can preference scores for discrete states be used to derive preference scores for an entire path of events? *Medical Decision Making* 17(1):42–55.

Lenert, L. A., Cher, D. J., Goldstein, M. K., Bergen, M. R., and Garber, A. 1998. The effect of search procedures on utility elicitations. *Medical Decision Making* 18(1):76–83.

Linden, G., Hanks, S., and Lesh, N. 1997. Interactive assessment of user preference models: The automated travel assistant. In *Proc. User Modeling '97*.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. San Francisco, CA: Morgan Kaufmann.

Poh, K. L., and Horvitz, E. 1993. Reasoning about the value of decision-model refinement: methods and application. In *Proc. UAI–93*, 174–182.

Shachter, R., and Kenley, C. 1989. Gaussian influence diagrams. *Management Science* 35:527–550.

von Neumann, J., and Morgenstern, O. 1947. *Theory of Games and Economic Behavior*. Princeton, N.J.: Princeton University Press, 2nd edition.