

Computational Approaches to Preference Elicitation

Darius Braziunas

Department of Computer Science
University of Toronto

2006

Contents

1	Introduction	1
2	Decision theory	2
2.1	Preferences under certainty	2
2.2	Preferences under uncertainty	3
2.3	Multiattribute outcomes	4
2.4	Additive utility representation	6
3	Main aspects of preference elicitation	7
3.1	“Classical” preference elicitation	7
3.1.1	Query types	8
3.1.2	Multiattribute elicitation	9
3.1.3	Problems with the classical paradigm	10
3.2	Decisions with partial preference information	10
3.2.1	Strict uncertainty	11
3.2.2	Bayesian uncertainty	13
3.3	Query selection criteria	13
3.3.1	Max regret reduction	14
3.3.2	Uncertainty reduction	14
3.3.3	Expected value of information	15
4	Research on preference elicitation	16
4.1	ISMAUT	16
4.2	Engineering design and configuration problems	19
4.3	Extensions of ISMAUT	20
4.4	Conjoint analysis	22
4.5	Analytic hierarchy process	24
4.6	Preference elicitation in AI	25
4.6.1	Minimax regret approach	25
4.6.2	Bayesian approach	27

1 Introduction

This is a survey of preference (or utility) elicitation from a computer scientist's perspective. Preference elicitation is viewed as a process of extracting information about user preferences to the extent necessary to make good or even optimal decisions. Devising effective elicitation strategies would facilitate building autonomous agents that can act on behalf of a user.

Artificial intelligence researchers have always been interested in developing intelligent decision aids with applications ranging from critical financial, medical, and logistics domains to low-stakes processes, such as product recommendation or automated software configuration. Decision theory provides solutions given the system dynamics and outcome utilities. However, user utilities are often unknown and vary more widely than decision dynamics. Since obtaining full preferences is usually infeasible, this presents a serious problem to the deployment of intelligent agents that make decisions or recommendations for users with distinct utilities. Therefore, preference elicitation emerges as one of the more important current challenges in artificial intelligence.

In this report, we consider both historical and current approaches to preference elicitation, concentrating on a few key aspects of the problem. We view preference elicitation and decision making as an inseparable sequential process. Although utility functions are hard to assess, partial information of user preferences might suffice to make good or optimal decisions. Preference elicitation can be driven to explore utility regions that are relevant for making decisions. On the other hand, knowledge of system dynamics and action constraints helps avoid eliciting useless utility information.

In the first part of the report, we describe "classical" decision analysis, consider ways to represent uncertainty over utility functions, and summarize various criteria for decision making with partial utility information. If uncertainty is too great to make good decisions, further preference elicitation is needed. One issue that arises in the elicitation process is which query to ask next. Intelligent querying strategies steer the elicitation process according to some agreed criteria and should be viewed as part of the combined preference elicitation and decision making process. We should note that our use of the term "query" also encompasses more implicit interactions with a user such as changing the presentation of a web page and observing the link followed by a user.

The second part of the report surveys research fields where preference elicitation plays a central role. Imprecisely specified multiattribute utility theory (ISMAUT) is one of the earlier attempts to consider decision making under partial preference information in classical decision analysis. Its extensions to engineering design and configuration problems have been influential in spurring recent interest in preference elicitation among artificial intelligence researchers. Conjoint analysis and analytical hierarchy process (AHP) methods were developed largely in isolation in the fields of marketing research and decision analysis; nonetheless, many issues involving preference elicitation are common. We finish by describing some of the recent advances in preference elicitation in AI.

2 Decision theory

"Theory of Games and Economic Behavior" by [von Neumann and Morgenstern, 1947]
 "The Foundations of Statistics" by [Savage, 1954]
 "Utility Theory for Decision Making" [Fishburn, 1970]
 "Decisions with Multiple Objectives: Preferences and Value Trade-offs" [Keeney and Raiffa, 1976]
 "The Foundations of Expected Utility" [Fishburn, 1982]
 "Decision Theory" [French, 1986]
 "Statistical decision theory" [French and Insua, 2000]

In this section we provide the background for decision-theoretic treatment of preferences. Decision theory lies at the intersection of many academic disciplines – statistics, economics, psychology, game theory, operations research, and others. Assuming a set of axioms for *rational* behavior, it provides a theory for modeling user preferences and making optimal decisions based on these preferences. The following summary of main concepts is based on [von Neumann and Morgenstern, 1947, Savage, 1954, Fishburn, 1970, Keeney and Raiffa, 1976, French, 1986, French and Insua, 2000].

In the basic formulation, a *decision maker* (DM) has to select a single alternative (or action) $a \in A$ from the set of available alternatives. An *outcome* (or *consequence*) $x \in X$ of the chosen action depends on the state of the world $\theta \in \Theta$. The consequence function $c : A \times \Theta \mapsto X$ maps each action and world state into an outcome. User preferences can be expressed by a *value*, or *utility*, function $v : X \mapsto \mathbb{R}$ that measures desirability of outcomes. The goal is to select an action $a \in A$ that leads to best outcomes. If the world state θ is known, the set of outcomes is equivalent to the set of alternatives; therefore, in such a case, we will often use these terms interchangeably. When uncertainty over world states is quantified probabilistically, *utility theory* prescribes an action that leads to the highest expected value.

The outcome space itself might be multidimensional. Most interesting problems fall in this category, and we survey some ways of exploiting the structure in multidimensional outcome spaces.

2.1 Preferences under certainty

We first consider decisions under certainty. The set of nature states now consists of a single state θ , and thus each action leads to a certain outcome. Preferences over outcomes completely determine the optimal action: a rational person would choose the action that results in the most preferred outcome.

Let X be a set of outcomes over which a preference relation is defined. The notation $x \succeq y$ means that a person *weakly* prefers outcome x to outcome y ; that is, outcome x is deemed to be as good as outcome y . The *weak preference* relation is commonly expected to satisfy the following two properties for the preferences to be considered *rational*:

$$\text{Comparability: } \forall x, y \in X, \quad x \succeq y \vee y \succeq x \tag{1}$$

$$\text{Transitivity: } \forall x, y, z \in X, \quad x \succeq y \wedge y \succeq z \implies x \succeq z \tag{2}$$

Weak preference is therefore a *total preorder* (or *weak order*) relation over the set of outcomes X . It is natural to think of weak preference as a combination of strict preference relation \succ and indifference relation \sim . The statement $x \succ y$ means that x is strictly preferred to y ; $x \sim y$ means that x is exactly as good as y . Formally, for any two elements $x, y \in X$

$$x \sim y \iff x \succeq y \wedge y \succeq x, \tag{3}$$

$$x \succ y \iff y \not\succeq x. \tag{4}$$

It follows that strict preference is a strict order (\succ is asymmetric and transitive), and indifference is an equivalence relation (\sim is reflexive, symmetric, and transitive).

Weak preferences can be represented compactly by a numerical function. An *ordinal value function* $v : X \mapsto \mathbb{R}$ represents or agrees with the ordering \succeq when for all $x, y \in X$

$$v(x) \geq v(y) \iff x \succeq y. \quad (5)$$

A *representation theorem* gives necessary and sufficient conditions under which some qualitative relation can be represented by a numerical ranking, or *scale*. In case of weak preferences, an agreeing ordinal value function can always be constructed if the outcome set X is finite or countably large. If X is uncountably large, then an agreeing ordinal value function exists if and only if X has a countable, order dense subset with respect to \succeq .¹

Ordinal value functions are unique up to *strictly increasing transformations*. Such functions are called *ordinal scale* functions. They contain only preference ranking information; thus, it would be meaningless to compare any linear combination of ordinal scale values (such as the average or difference of outcome values).

2.2 Preferences under uncertainty

In many settings, the consequences of an action are uncertain. Modern utility theory is based upon the fundamental work of von Neumann and Morgenstern [1947]. In this theory, uncertainty is quantified probabilistically, and a rational decision maker is capable of expressing preferences between *lotteries*, or probability distributions over a *finite* set of outcomes.

A *simple lottery*, where outcome x_i is realized with probability p_i , is conventionally denoted as

$$l = \langle p_1, x_1; p_2, x_2; \dots; p_n, x_n \rangle. \quad (6)$$

It is common to omit outcomes with zero probabilities in the lottery notation. When the lottery contains only two outcomes with positive probabilities, it will sometimes be abbreviated as $\langle p, x; 1 - p, x' \rangle \equiv \langle x, p, x' \rangle$.

A more general type of lottery is a *compound* lottery, where the outcomes themselves are simple lotteries: $l' = \langle p_1, l'_1; p_2, l'_2; \dots; p_n, l'_k \rangle$. Any compound lottery can be *reduced* to an equivalent simple lottery where the final outcomes are realized with same probabilities. Using the vector notation, the reduced lottery obtained from l' is simply $p_1 l'_1 + p_2 l'_2 + \dots + p_n l'_k$. An important assumption about preferences over lotteries is that the decision maker views any compound lottery and its reduction as equivalent; that is, only the ultimate probabilities of outcomes matter. It therefore suffices to consider preferences over the set of simple lotteries.

As in the case of certainty, the rational decision maker is assumed to have a complete and transitive preference ranking \succeq over the set of simple lotteries L . The *continuity*, or *Archimedean*, axiom states that no alternative is infinitely better (or worse) than others:

$$\begin{aligned} \text{Continuity: } \quad & \forall l_1, l_2, l_3 \in L, \\ & l_1 \succ l_2 \succ l_3 \implies \langle l_1, p, l_3 \rangle \succ l_2 \succ \langle l_1, q, l_3 \rangle, \\ & \text{for some } p, q \in (0, 1). \end{aligned} \quad (7)$$

The continuity axiom is required for existence of a *utility function* $u : L \mapsto \mathbb{R}$ that represents the preference relation \succeq on simple lotteries. An additional *independence* axiom is necessary to impose a very convenient

¹ $Y \subseteq X$ is *order dense* with respect to \succeq if $\forall x, z \in X$ such that $x \succeq z$, there exists $y \in Y$ such that $x \succeq y \succeq z$. For example, rational numbers form a countable, order dense subset of real numbers with respect to the order \geq .

linear structure on the utility function $u(\cdot)$:

$$\begin{aligned} \text{Independence: } \quad & \forall l_1, l_2, l_3 \in L, \text{ and } p \in (0, 1), \\ & l_1 \succ l_2 \implies \langle l_1, p, l_3 \rangle \succ \langle l_2, p, l_3 \rangle. \end{aligned} \quad (8)$$

Independence axiom requires that preferences over l_1 and l_2 carry over to compound lotteries involving some other lottery l_3 .

The most important result that follows is the *expected utility* representation theorem. It states that if and only if the weak preference relation on simple lotteries is (1) complete, (2) transitive, (3) satisfies continuity axiom, and (4) satisfies independence axiom, then there exists an *expected* or *linear* utility function $u : L \mapsto \mathbb{R}$ which represents \succeq . A utility function $u(\cdot)$ has the following properties:

$$\begin{aligned} (1) \quad & u(l) \geq u(l') \iff l \succeq l', \\ (2) \quad & u(\langle l, p, l' \rangle) = pu(l) + (1-p)u(l'), \quad \forall l, l' \in L, \text{ and } p \in [0, 1]. \end{aligned} \quad (9)$$

We can identify any outcome $x \in X$ with a degenerate lottery $l^x = \langle 1, x; 0, \dots \rangle$, where outcome x occurs with certainty. This allows us to extend the preference relation \succeq on simple lotteries to outcomes. The utility of outcome x is then the same as that of the corresponding degenerate lottery: $u(x) = u(l^x)$. Using induction and linearity of the utility function $u(\cdot)$, it can be shown that the utility of any simple lottery $l = \langle p_1, x_1; p_2, x_2; \dots; p_n, x_n \rangle$ is the expected value of its outcomes.

$$u(l) = u(\langle p_1, x_1; p_2, x_2; \dots; p_n, x_n \rangle) = \sum_{i=1}^n p_i u(x_i). \quad (10)$$

This key result allows us to represent preferences over an infinite set of simple lotteries by a utility function over a finite set of outcomes.

2.3 Multiattribute outcomes

In practice, the set of outcomes X is often endowed with multidimensional structure. For example, each alternative in A can be evaluated on several criteria, or *attributes*. Under certainty, action $a \in A$ maps to a point in a *multiattribute* space; under uncertainty, it maps to a distribution over points in that space. The goal of *multiattribute utility theory* (MAUT) is to investigate numerical representations that reflect structure in user preferences over multiattribute spaces. [Keeney and Raiffa, 1976] remains the main reference for MAUT.

Assume a set of attributes X_1, X_2, \dots, X_n . Each attribute is either a finite or infinite set of possible levels, or values (attributes can be also thought of as variables; for ease of notation we use X_i to refer to its domain as well). The set of all outcomes $\mathbf{X} = X_1 \times \dots \times X_n$ is the Cartesian product of attribute levels. Given an index set $I \subseteq \{1, \dots, n\}$, we define $\mathbf{X}_I = \times_{i \in I} X_i$ to be the set of *partial outcomes* restricted to attributes in I , and \mathbf{x}_I to be the same restriction of a specific outcome \mathbf{x} . I^C denotes I 's complement.

If preferences over multiattribute outcomes exhibit sufficient structure, a preference relation can be modeled more concisely, and utility functions can be decomposed into *subutility* functions, defined over subsets of attributes. The simplest *independence condition* is called *preferential independence*. Attributes in I are preferentially independent of the remaining attributes if

$$\begin{aligned} & (\mathbf{x}_I, \mathbf{y}) \succeq (\mathbf{x}'_I, \mathbf{y}) \text{ for some } \mathbf{y} \in \mathbf{X}_{I^C} \\ \implies & (\mathbf{x}_I, \mathbf{y}') \succeq (\mathbf{x}'_I, \mathbf{y}') \text{ for all } \mathbf{y}' \in \mathbf{X}_{I^C}. \end{aligned} \quad (11)$$

That is, as long as attributes not in I are fixed to some level \mathbf{y} , the preferences over \mathbf{X}_I do not depend on the setting of remaining attributes. Therefore, a statement $\mathbf{x}_I \succeq \mathbf{x}'_I$, *ceteris paribus* (all else being equal), is a concise way of stating $(\mathbf{x}_I, \mathbf{y}) \succeq (\mathbf{x}'_I, \mathbf{y})$ for all $\mathbf{y} \in \mathbf{X}_{I^c}$.

Ceteris paribus preferential statements provide a natural language for expressing multiattribute preferences. In AI, *ceteris paribus* assumptions are central to some of the *qualitative decision theories* [Doyle and Thomason, 1999]. The *logic of relative desire*, introduced by Doyle et al. [1991], as well as later work by Doyle and Wellman [1994] interpret planning goals as qualitative preference statements over models using all-else-being-equal semantics. In a more recent article, McGeachie and Doyle [2004] provide algorithms for computing *ordinal* value functions based on qualitative *ceteris paribus* preference statements.

One issue not addressed by *ceteris paribus* logic theories was compact and efficient *representation* of preferential independence statements. A *CP-net*, introduced by Boutilier et al. [1999], is a popular graphical model that exploits conditional preferential independence among attributes. To create the structure of a CP-net, for each attribute X_i , a user must indicate which other attributes — *parents* of X_i — impact the preferences over values of attribute X_i . Then, for each possible instantiation of the parents of X_i , the user provides a qualitative preference relation over the values of X_i , all else being equal. Given a CP-net, the *ceteris paribus* semantics induces a *partial* order over full outcomes. Besides providing a compact and natural representation of preferences, a CP-net N can be used to perform preferential comparison between full outcomes (“Does N entail $\mathbf{x} \succ \mathbf{x}'$?”), partial outcome optimization (“What is the best outcome \mathbf{x} given N ?”), and outcome ordering (“Is there some ranking in which $\mathbf{x} \succ \mathbf{x}'$?”).

While partial outcome optimization and outcome ordering are computationally tractable (polynomial in the size of the network), dominance testing is more complicated. In general, answering dominance queries is PSPACE-complete; however, polynomial algorithms exist for tree and polytree structured networks. Another complication is cyclicity: while quite natural in certain settings, cyclical networks are not guaranteed to have a satisfiable preference ranking. Because satisfiability testing can be hard, most research on CP-nets (including complexity results mentioned before) is limited to acyclic networks. Finally, we should note that introduction of *ceteris paribus indifference* statements can also lead to unsatisfiable networks.

CP-nets have been extended to deal with hard [Boutilier et al., 2004a] and soft [Prestwich et al., 2005] constraints. Another interesting generalization is the TCP-net, which adds conditional importance relations among variables [Brafman and Domshlak, 2002].

The idea of preferential independence extends to preferences over lotteries, leading to the notion of *utility independence*. Let L be the set of all simple lotteries on \mathbf{X} , and L_I — the set of all lotteries on \mathbf{X}_I . For $l \in L$, l_I is the marginal of l on \mathbf{X}_I . User preferences for \mathbf{X}_I are *utility independent* of \mathbf{X}_{I^c} if preferences over marginal lotteries over \mathbf{X}_I , when the levels of \mathbf{X}_{I^c} are fixed to \mathbf{y} , do not depend on that fixed level \mathbf{y} . A graphical *UCP-net* model [Boutilier et al., 2001] that exploits conditional utility independence among attributes is the quantitative analogue of CP-nets.

Preference independence and utility independence always involve a relationship between two complementary sets of attributes. Independencies between arbitrary subsets of attributes will be discussed later, when describing conditions for additive utility representation. Finally, we should mention the notion of Pareto-optimal alternatives when discussing multiattribute outcomes. Consider a set of multiattribute alternatives A , such that each attribute is preferentially independent of the remaining attributes. Without loss of generality, we can assume that preferences are monotonically increasing with the value of each attribute. Then, alternative \mathbf{x} *dominates* \mathbf{x}' if

$$x_i \succeq x'_i \text{ for } i = 1, \dots, n \text{ with } x_i \succ x'_i \text{ for at least one } i. \quad (12)$$

The *Pareto optimal set* (also known as *efficient set* or *admissible set*) is the set of all nondominated alterna-

tives in A . It is common to restrict the set alternatives to the Pareto optimal set, because a Pareto-dominated alternative cannot be optimal.

2.4 Additive utility representation

When utility or value functions have additive representations, many techniques of preference elicitation are similar in both certain and uncertain settings. The term “utility function” will be used as a synonym for a value function, unless explicitly noted otherwise.

Since the number of outcomes is exponential in the number of attributes, specifying the utility value for each outcome is infeasible in many practical applications. However, $u(\cdot)$ can be expressed concisely if it exhibits sufficient structure. *Additive independence* [Keeney and Raiffa, 1976] is one structural assumption commonly used in practice. Under certainty, additive independence requires that *all* subsets of attributes be mutually preferentially independent of their compliments. Under uncertainty, the decision maker has to be indifferent among lotteries that have same marginals on each attribute. When additive independence holds, $u(\cdot)$ can be written as a sum of single-attribute *subutility functions*:

$$u(\mathbf{x}) = \sum_{i=1}^n u_i(x_i) = \sum_{i=1}^n w_i v_i(x_i). \quad (13)$$

This simple factorization exploits subutility functions $u_i(x_i) = w_i v_i(x_i)$, which can be written as product of *local value functions* v_i and *scaling factors*, or *weights*, w_i . The two representations — the sum of attribute subutility functions and the sum of weighted local value functions — are equivalent; the weighted representation is commonly used under the assumption that weights form a simplex (i.e., $\sum_i w_i = 1$, $w_i \geq 0$) and local value functions are normalized to be in the range $[0;1]$.

If attributes x_i are numerical and a value function can be written as

$$u(\mathbf{x}) = \sum_{i=1}^n \lambda_i x_i, \quad (14)$$

then it is *linear*. Such functions are quite commonly assumed in operations research, cost-benefit analysis, and economics. In addition to conditions required for existence of additive value functions, there is an additional property of *constant relative tradeoff* between every pair of attributes that has to be satisfied. A pair of attributes i and j has a constant relative tradeoff ρ_{ij} if the decision maker is always indifferent between some outcome \mathbf{x} and an outcome obtained by increasing x_i and decreasing x_j in the ratio $\rho_{ij} : 1$. Linear functions are therefore measured on a more restrictive *ratio* scale: they are unique up to scaling by a positive constant.

While additive models are by far the most commonly used in practice, *generalized additive independence* (GAI) models have recently gained attention because of their additional flexibility (see, e.g., [Bacchus and Grove, 1995, Boutilier et al., 2001, 2003b, Gonzales and Perny, 2004, Boutilier et al., 2005, Braziunas and Boutilier, 2005]). The conditions under which a GAI model provides an accurate representation of a utility function were defined by Fishburn [1967a, 1970], who introduced the model.² GAI is a generalization of the additive model, where independence holds among certain *subsets* of attributes, rather than individual attributes.

²Fishburn used the term *interdependent value additivity*; Bacchus and Grove [1995] dubbed the same concept GAI, which seems to be more commonly used in the AI literature.

Let $\{I_1, \dots, I_m\}$ be a collection of nonempty subsets of $\{1, \dots, n\}$. Also, recall that l_I denotes the marginal of the lottery l on the attributes in I . The sets of attributes indexed by I_1, \dots, I_m are (*generalized*) *additively independent* if and only if

$$[(l_{I_1}, \dots, l_{I_m}) = (l'_{I_1}, \dots, l'_{I_m})] \implies l \sim l', \quad (15)$$

that is, if and only if the decision maker is indifferent between two lotteries whenever their marginal distributions on $\mathbf{X}_{I_1}, \dots, \mathbf{X}_{I_m}$ are the same. When generalized additive independence holds, the utility of a multiattribute outcome can be written as a sum of subutilities involving GAI subsets of attributes:

$$u(\mathbf{x}) = \sum_{i=1}^m u_i(\mathbf{x}_{I_i}). \quad (16)$$

3 Main aspects of preference elicitation

The increased interest in automated decision support tools in recent years has brought the problem of *automated preference elicitation* to the forefront of research in decision analysis [Dyer, 1972, White et al., 1984, Salo and Hämäläinen, 2001] and AI [Chajewska et al., 1998, 2000, Boutilier, 2002]. The goal of automated preference elicitation is to devise algorithmic techniques that will guide a user through an appropriate sequence of queries or interactions and determine enough preference information to make a good or optimal decision.

In this section, we concentrate on a few key aspects of the preference elicitation problem. The previous section dealt with various complete representations of preference information. Here, we are interested in the actual process of acquiring such information as well as making decisions with partially elicited utility functions. Therefore, we address issues of how to represent uncertainty over possible utility functions, how to make decisions without full knowledge of user preferences, and how to intelligently guide the elicitation process by taking into account the cost of interaction and potential improvement of decision quality.

3.1 “Classical” preference elicitation

Preference (or utility) elicitation is a process of assessing preferences or, more specifically, utility functions. Utility elicitation literature is as old as utility theory itself; first attempts to describe procedures for evaluating utility functions date back to the 1950s. In the “classical” setting, a *decision analyst’s* task is to help elicit a *decision maker’s* preferences. Once those preferences are extracted, the decision analyst calculates an optimal course of action according to the utility theory, and *recommends* it to the decision maker [Keeney and Raiffa, 1976, Howard and Matheson, 1984, French, 1986].

There are many techniques for evaluating utility functions, and the whole process of elicitation is “as much of an art as it is a science” [Keeney and Raiffa, 1976]. A classical approach, involving an interaction between the decision analyst and the decision maker, usually consists of five steps [Keeney and Raiffa, 1976, Farquhar, 1984]. During the *preparation for assessment*, the DM is acquainted with the decision problem, possible outcomes or attributes, and various aspects of the elicitation procedure. The next stage is *identification of relevant qualitative characteristics* of DM’s preferences. This could include determining the properties of the utility function (such as continuity or monotonicity in case of numerical attributes), best and worst outcomes or attribute levels, and independence relations among attributes for structured outcome spaces. The central part of the procedure is *specification of quantitative restrictions* and *selection of a utility function*. Here, the decision analyst asks various queries, some of which are described below, in an attempt to model DM’s preferences by a *completely* specified utility function. Most of the approaches described in this survey depart from the classical form of elicitation because of the complexity of this

task. The last step usually involves *checks for consistency* and *sensitivity analysis*. When inconsistencies are detected, the DM is asked to revise her preferences. The goal of sensitivity analysis is to check the sensitivity of the output (which, in most cases, is the decision recommended by the decision analyst) to the inputs — the utility model and DM’s responses.

3.1.1 Query types

The nature of queries is an integral part of the preference elicitation problem. Some queries are easy to answer, but do not provide much information; and, vice versa, informative queries are often costly. Another tradeoff to consider is the complexity of selecting the right query versus its potential usefulness. Such aspects of preference elicitation depend on *query types*.

We survey some queries that are commonly used in decision analysis and describe their main characteristics. “Global” queries are applicable to situations where either the set of outcomes does not have any structure, or, in case of multiattribute problems, that structure is ignored and only full, or *global*, outcomes are considered. In most multiattribute problems, people can meaningfully compare outcomes with no more than five or six attributes [Green and Srinivasan, 1978]. Therefore, most of the global queries have “local” counterparts that apply to a subset of attributes.

We assume that preferences over the set of outcomes X can be expressed by a utility function $u(\cdot)$, and consider queries that help assess this function. Utility functions are unique up to positive affine transformations; therefore, if $u(\cdot)$ represents the preference relation \succeq , then so does $u'(\cdot) = au(\cdot) + b$, with $a \in \mathbb{R}^+$, $b \in \mathbb{R}$. Without loss of generality, we can therefore normalize any utility function to lie between 0 and 1.

Order comparison *Order queries* are very simple queries that ask the user to compare a pair of alternatives x and y ; the user might prefer x to y , y to x , or be indifferent between the two. Such queries are very common in practice (they are central in ISMAUT and conjoint analysis, for example) and usually require little cognitive effort from the user. Unfortunately, often they are not too informative.

More complicated comparison queries ask the user to pick the most preferred alternative from the set of k alternatives. This rather easy task actually provides $k - 1$ preference relations (the selected alternative is preferred to all remaining choices), and is widely used in choice-based conjoint analysis. At the most extreme, a *total ranking* query expects the user to rank all specified alternatives; answering such a query would provide preference information relating every pair of alternatives.

Most other utility elicitation queries involve degenerate lotteries, or *gambles*, with only two outcomes. As before, we use abbreviated notation $\langle x, p, x' \rangle$ for a lottery where x occurs with probability p and x' occurs with probability $1-p$. We consider a general query expression $\langle x, p, x' \rangle \gtrsim y$, where everything except one item is specified, and the user is asked to provide the value of the item that would make the expression true. In the expression, x, x', y are outcomes in X , p is a probability, and \gtrsim is either \succeq or \prec ³. The following terminology and classification follows [Farquhar, 1984], who describes queries for all possible combinations of known and unknown quantities in the query expression, as well as more general queries involving two gambles.

Probability equivalence Probability equivalence queries elicit an indifference probability p for which $\langle x, p, x' \rangle \sim y$. In a standard gamble case, when $x = x^\top$ and $x' = x^\perp$, the query simply asks to specify the utility of y , and is therefore sometimes called a *direct utility query*. While such queries have been used in research papers [Keeney and Raiffa, 1976, Gonzales and Perny, 2004], it is unlikely that users can provide

³Or, more generally, one of the three preference relations \succ, \sim and \prec .

exact utility values for outcomes in real-world situations. One possible generalization is to ask for *bounds* on the utility value; these bounds can then be narrowed by asking binary comparison queries described below [Boutillier et al., 2003b].

Preference comparison In a preference comparison between a gamble $\langle x, p, x' \rangle$ and a sure outcome y , a user is asked to specify a relation (\succeq or \prec) that holds between the two. When the gamble is a standard gamble, $u(\langle x^\top, p, x^\perp \rangle) = p$, and the query becomes equivalent to “Is $u(y) \geq p$?” with possible $\{yes, no\}$ responses.⁴ Such query is called a *standard gamble comparison query*.

Standard gamble comparison queries are common in classical decision analysis literature [Keeney and Raiffa, 1976]. More recently, such queries have been used by Chajewska and Koller [2000], Boutillier [2002], Wang and Boutillier [2003], Boutillier et al. [2003b, 2005], Braziunas and Boutillier [2005], and others.

Implicit queries Until now, we assumed that a query is an explicit question posed by the decision support system, and a response is a user’s reaction to the query. However, queries and responses can be much more general. The system could pose “implicit” queries by changing the user environment (such as options available on the web page), and observing the user’s behavior (links followed, time spent on the page, etc.). Or, the user can be asked to view a fragment of some action policy, and asked to critique the actions. A related theoretical framework is *inverse reinforcement learning* [Ng and Russell, 2000, Chajewska et al., 2001]. The goal is to recover the reward function (preferences) of an agent by observing execution of an optimal policy.

The concept of *revealed preference* in economics [Mas-Colell et al., 1995] is also related to the topic of implicit queries. Here, the emphasis is on descriptive, rather than prescriptive, aspects of human decision making. Observable choices that people make faced with an economic decision provide the primary basis for modeling their behavior. A preference relation does not exist *a priori*, but could be *derived* (or revealed) given observed choices that follow certain axioms of rationality.

3.1.2 Multiattribute elicitation

All the techniques described above could be also employed to simplify utility elicitation of structured outcomes. To illustrate the main concepts, we consider the case of eliciting an additive utility function

$$u(\mathbf{x}) = \sum_{i=1}^n u_i(x_i) = \sum_{i=1}^n w_i v_i(x_i), \quad (17)$$

where $u_i(x_i)$ are subutility functions that can be written as a product of *local value functions* v_i and *scaling factors*, or *weights*, w_i .

The assumed utility independence among attributes allows elicitation to proceed *locally*: specifically, each $v_i(\cdot)$ can be elicited independently of other attribute values using any of the techniques described above. Since attributes are preferentially independent, each attribute’s best and worst levels (we shall call them *anchor levels*) can be determined separately. Let x_i^\top and x_i^\perp denote the best and worst levels, respectively, of attribute i . Local value functions $v_i(\cdot)$ can be determined by locally measuring values of attribute levels with respect to the two anchor levels. What remains is to bring all the local value scales to the common global utility scale. Essentially, we need to find the true utility of all “anchor” outcomes x_i^\top and x_i^\perp relative to some default outcome \mathbf{x}^0 (it is customary to choose the worst outcome as default outcome, and set its utility to 0). Then, eliciting $u(x_i^\top, \mathbf{x}_{iC}^0) = u_i(x_i^\top)$ and $u(x_i^\perp, \mathbf{x}_{iC}^0) = u_i(x_i^\perp)$ for

⁴We should note that despite the equivalence, comparing an outcome y and a standard lottery $\langle x, p, x' \rangle$ might be psychologically easier than providing a response regarding the utility value of y .

all attributes would ensure consistent scaling of subutility functions. Because additive utility functions are unique up to positive affine transformations, it is usually assumed that both the global utility functions and local value functions are scaled to lie between 0 and 1; the weights are also normalized such that $w_i \geq 0$ and $\sum w_i = 1$. In a normalized additive utility function, scaling factors w_i , which reflect attribute contributions to the overall utility function, are simply equal to $u_i(x_i^\top)$.

The scaling factors can be determined by asking utility queries involving full outcome lotteries. The simplest way to elicit w_i is to find the utility of $(x_i^\top, \mathbf{x}_{iC}^\perp)$ for all i :

$$w_i = u(x_i^\top, \mathbf{x}_{iC}^\perp). \quad (18)$$

More generally, we need $n - 1$ independent linear equations involving the unknown scaling constants. Such equations could be chosen in a manner that reduces the cognitive burden of the DM (for example, by carefully choosing queries that do not involve extreme attribute values).

More general multiattribute utility functions, such as multiplicative or GAI, can also be elicited using ideas of local value function elicitation and global scaling [Fishburn, 1967b, Keeney and Raiffa, 1976, Fishburn, 1977, Braziunas and Boutilier, 2005].

3.1.3 Problems with the classical paradigm

Complete preference information is often unattainable in practice. In many realistic domains where the outcome space is large, it is unreasonable to expect a user to provide preference information about every outcome. In multiattribute settings with more than, say, ten attributes, complete preference elicitation becomes virtually impossible, as the number of alternatives is exponential in the number of attributes.

Elicitation of quantitative utilities brings additional difficulties. Queries involving numbers and probabilities are cognitively hard to answer; most users are not experts and therefore require preliminary training. Real case studies often provide evidence of inconsistent responses, errors, and various forms of biases. Eliciting preferences might be costly, too; costs can be cognitive (hours of human effort in answering questionnaires), computational (calculating a value of certain alternative might involve solving complicated optimization problems or running simulations), financial (hiring a team of experts to analyze potential business strategies), and others.

Furthermore, from the AI perspective, preference elicitation presents a “bottleneck” for designing automated decision aids ranging from critical financial, medical, and logistics domains to low-stakes processes, such as product recommendation or automated software configuration. For making optimal decisions, we need to know both the decision dynamics and outcome utilities. In many situations, the dynamics is known (elicitation and representation of complex probability models is a well-researched area of AI). However, user preferences are often unknown, and, furthermore, they vary considerably from user to user (while the system dynamics is often fixed for all users). Designing effective preference elicitation techniques is therefore an important problem facing AI.

When costs of elicitation are taken in to account, it becomes clear that decisions might have to be made with partial preference information, if elicitation costs start to exceed potential improvement of decisions. Viewing utility elicitation as an integral part of the decision process is a promising paradigm for tackling the preference elicitation problem.

3.2 Decisions with partial preference information

If the utility function is not fully known and further elicitation not possible, what criteria should be used for making good decisions with available information? It turns out that criteria proposed for dealing with state uncertainty in classical decision theory (such as maximum expected utility, minimax regret, maximin) can be applied to situations where utility functions themselves are uncertain. The analogy extends to

both common representations of utility function uncertainty: Bayesian, where we can keep track of the probability distribution over possible utility functions, and strict uncertainty, defined by the set of *feasible* utility functions.

3.2.1 Strict uncertainty

“UCP-Networks: A Directed Graphical Representation of Conditional Utilities” [Boutilier et al., 2001]
 “Preference Ratios in Multiattribute Evaluation (PRIME)” [Salo and Hämäläinen, 2001]
 “Preference programming” [Salo and Hämäläinen, 2004]
 “Cooperative Negotiation in Autonomic Systems using Incremental Utility Elicitation” [Boutilier et al., 2003a]
 “Incremental Utility Elicitation with the Minimax Regret Decision Criterion” [Wang and Boutilier, 2003]

Under strict uncertainty, knowledge about a user’s utility function is characterized by the feasible utility set U . This set is updated (reduced) when relevant preference information is received during an elicitation process. The following is a non-exhaustive list of decision criteria that could be used for making decisions with partial utility information under strict uncertainty. The set of outcomes is X , and the goal is to choose the best alternative x^* when the set of feasible utility functions is U .

Maximin return Without distributional information about the set of possible utility functions U , it might seem reasonable to select an outcome whose worst-case return is highest:

$$x^* = \arg \max_{x \in X} \min_{u \in U} u(x). \quad (19)$$

Maximin decision is sometimes called *robust* because it provides an *ex post* security guarantee. Maximin was proposed by Wald [1950], and mentioned by Salo and Hämäläinen [2004] for the case of uncertain utilities.

Hurwicz’s optimism-pessimism index Maximin return is a pessimistic criterion, because the decision maker prepares for the worst realization of the utility function. *Maximax return* criterion is the optimistic counterpart to maximin. Supposing that maximin and maximax are too extreme, Hurwicz proposed to use a weighted combination of the minimum and maximum possible values [French, 1986]. For the case of strict uncertainty over utility functions, this criterion would choose

$$x^* = \arg \max_{x \in X} \left[\alpha \min_{u \in U} u(x) + (1 - \alpha) \max_{u \in U} u(x) \right], \quad (20)$$

where α is the *optimism-pessimism index* of the decision maker. Hurwicz’s optimism-pessimism criterion generalizes minimax and maximax, as well as the *central values* criterion favored by Salo and Hämäläinen [2001]. Central values rule prescribes an outcome whose mid-point of the feasible utility interval is largest, which is equivalent to setting the optimism-pessimism index α to 0.5.

Minimax regret Minimax regret criterion was first described by Savage [1951] in the context of uncertainty over world states, and advocated by Boutilier et al. [2001] and Salo and Hämäläinen [2001] for robust decision making with uncertain utility functions. The main idea is to compare decisions for *each* state of uncertainty. The *maximum regret* of choosing x is $MR(x, U) = \max_u \max_{x'} [u(x') - u(x)]$. Minimax regret optimal decision minimizes the worst-case loss with respect to possible realizations of the utility function:

$$x^* = \arg \min_{x \in X} MR(x, U) = \arg \min_{x \in X} \max_{u \in U} \max_{x' \in X} [u(x') - u(x)]. \quad (21)$$

Various applications of decision making with minimax regret criterion have been researched by Boutilier et al. [2001, 2003a,b], Wang and Boutilier [2003], Boutilier et al. [2004b, 2005], Patrascu et al. [2005].

The largest drawback of this criterion is the failure to satisfy the *principle of independence of irrelevant alternatives*. According to this principle, the ranking between two alternatives should be independent of other available alternatives (for example, the violation of this principle could result in a situation where $x \succeq y$ if option z is available, and $y \succeq x$ otherwise). We should note, however, that the principle of independence of irrelevant alternatives is by no means universally accepted as a prerequisite for rational decision making.

Principle of Insufficient Reason This criterion dates back to Pierre Laplace and Jacob Bernoulli, who maintained that complete lack of knowledge about the likelihood of world states should be equivalent to all states having equal probability. Therefore, following this *principle of insufficient reason*, an optimal decision maximizes the mean value of possible outcomes:

$$x^* = \arg \max_{x \in X} E_{u \in U}^{\pi} [u(x)], \quad (22)$$

where the π is the uniform distribution over U . This criterion is mentioned by Salo and Hämäläinen [2001] as *central weights* decision rule, and is implicitly employed by Iyengar et al. [2001], Ghosh and Kalagnanam [2003], Toubia et al. [2004], where uncertainty over additive utility functions is characterized by linear constraints on attribute weights. The “center” of the weight polytope could be its actual mass center (i.e., the mean under uniform distribution), or some approximation thereof — a point that minimizes maximal distance to constraint hyperplanes, the center of a bounding ellipsoid, or the average of uniformly sampled points from inside the region.

Acceptability index Finally, there are methods that recommend choosing an alternative based on the set size of supporting utility functions. Lahdelma et al. [1998] introduce *stochastic multiobjective acceptability analysis* (SMAA), which applies to settings where uncertainty over additive utility functions can be described by linear constraints on the $n - 1$ dimensional weight simplex W . Each alternative is associated with a region of W in which it is optimal. Alternatives are ranked according to *acceptability index*, which is the normalized volume of the weight region in which it is optimal. An alternative with the highest acceptability index is in some sense most likely to be optimal. Like the minimax regret criterion, the acceptability index criterion does not satisfy the principle of independence of irrelevant alternatives.

Various criteria for decision making under strict uncertainty can be grouped into categories based on their general properties. Maximax, maximin, and central values (i.e., optimism-pessimism index) are based on extreme possible values of outcomes. Center-based criteria pick a “representative” point in the space of feasible utilities. Finally, minimax regret is qualitatively different from the other criteria because it considers pairwise value differences between outcomes.

Unfortunately, different decision rules might prescribe different alternatives. The choice of a decision rule under strict uncertainty should be carefully considered by the decision maker before the elicitation process. French [1986] provides an extensive discussion and critique of various decision criteria under strict uncertainty.

3.2.2 Bayesian uncertainty

“Adaptive Utility” [Cyert and de Groot, 1979]
 “Decision Making with an Uncertain Utility Function” [de Groot, 1983].
 ”Utilities as Random Variables: Density Estimation and Structure Discovery” [Chajewska and Koller, 2000]
 ”Making Rational Decisions Using Adaptive Utility Elicitation” [Chajewska et al., 2000]
 ”A POMDP Formulation of Preference Elicitation Problems” [Boutilier, 2002]
 “On the Foundations of *Expected* Expected Utility” [Boutilier, 2003]
 “Local Utility Elicitation in GAI Models” [Braziunas and Boutilier, 2005]

A true Bayesian would likely reject the very notion of strict uncertainty. An optimal decision is simply the one that maximizes expected value, where expectation is taken with respect to a prior probability distribution π over the set of feasible utilities U :⁵

$$x^* = \arg \max_{x \in X} E_{u \in U}^{\pi} [u(x)] = \arg \max_{x \in X} EU(x, \pi), \quad (23)$$

where $EU(x, \pi)$ is the expected utility of outcome x when π is the probability distribution over utilities. In the case of additional uncertainty over world states, the goal is to maximize *expected* expected utility [Boutilier, 2003].

While most recent work on decision making using distributions over utility functions has been done within the AI community [Chajewska and Koller, 2000, Chajewska et al., 2000, Boutilier, 2002, Braziunas and Boutilier, 2005], the origins of this approach can be traced back to much earlier research in game theory and decision theory. Cyert and de Groot [1979] and de Groot [1983] propose the concept of *adaptive utility*, where a decision maker does not fully know her own utility function until a decision is made. Uncertainty is quantified as a probability distribution over utility function parameters. The distribution is updated by comparing expected utility of an outcome versus its actual utility, which becomes known after the decision is made. Weber [1987] also discusses using expectations over utility functions as a possible criterion for decision making with incomplete preference information. In a related context, probabilistic modeling of possible payoff functions provides the foundation to the well-established field of Bayesian games [Harsanyi, 1967, 1968].

Boutilier [2003] investigates the conditions under which it is reasonable to model uncertainty over functions measured on the interval scale. By appealing to the foundational axioms of utility theory, it can be shown that the functions are required to be *extremum equivalent*, i.e., they have to share the same best and worst outcomes.

An important issue in the Bayesian approach to modeling uncertainty over utility functions is the choice of prior probability distributions. Ideally, the probability model would be closed under updates (otherwise, it needs to be refit after each response) and flexible enough to model arbitrary prior beliefs. Mixtures of Gaussians [Chajewska et al., 2000], mixtures of truncated Gaussians [Boutilier, 2002], mixtures of uniforms [Boutilier, 2002, Wang and Boutilier, 2003, Braziunas and Boutilier, 2005], and Beta distributions Abbas [2004] are among possibilities proposed in the literature. Priors can also be learned from data — Chajewska et al. [1998] describe a way to cluster utility functions using a database of utilities from a medical domain.

3.3 Query selection criteria

A central issue in preference elicitation is the problem of which query to ask at each stage of the process. The value of a query is generally determined by combining the values of possible situations resulting from

⁵With a prior over utilities, the same decision also minimizes expected regret.

user responses. Similar to decision making with incomplete information, query selection is driven by the ultimate goals of the decision support system. Query selection criteria include fastest reduction of minimax regret or uncertainty, or achieving optimal tradeoff between elicitation costs and predicted improvement in decision quality.

3.3.1 Max regret reduction

"Incremental Utility Elicitation with the Minimax Regret Decision Criterion" [Wang and Boutilier, 2003]
"Cooperative Negotiation in Autonomic Systems using Incremental Utility Elicitation" [Boutilier et al., 2003a]
"New Approaches to Optimization and Utility Elicitation in Autonomic Computing" [Patrascu et al., 2005]
"Eliciting Bid Taker Non-price Preferences in (Combinatorial) Auctions" [Boutilier et al., 2004b]
"Constraint-based Optimization with the Minimax Decision Criterion" [Boutilier et al., 2003b]
"Regret-based Utility Elicitation in Constraint-based Decision Problems" [Boutilier et al., 2005]

The minimax regret decision criterion provides bounds on the quality of the decision made under strict uncertainty. When the potential regret associated with each decision is too high, more utility information needs to be elicited. A decision support system can query the user until the minimax regret reaches some acceptable level, elicitation costs become too high, or some other termination criterion is met.

Each possible response to a utility query results in a new decision situation with a new level of minimax regret (the level of regret cannot increase with more information). The problem is to estimate the value of a query based on the value of possible responses. For example, one could select the query with the best worst-case response, or the query with the maximum average or expected improvement [Wang and Boutilier, 2003].

Minimax regret reduction queries are also used in the autonomic computing scenario [Boutilier et al., 2003a, Patrascu et al., 2005], eliciting values of non-price features in combinatorial auctions [Boutilier et al., 2004b], and optimizing constrained configurations [Boutilier et al., 2003b, 2005]. A more detailed description of these methods is postponed till Section 4.6.1.

3.3.2 Uncertainty reduction

There is a variety of methods from diverse research areas, such as conjoint analysis and ISMAUT (see Sections 4.3 and 4.4), whose central idea is to choose queries that reduce the uncertainty over utility functions as much as possible. The set of possible utility functions is commonly represented as a convex polytope in the space of utility function parameters. Each query bisects the polytope by adding a linear constraint. Since the responses are not known beforehand, various heuristics are used to choose the next query. Such heuristics consider the size parity of volumes [Iyengar et al., 2001], as well as their shape [Ghosh and Kalagnanam, 2003, Toubia et al., 2004].

Abbas [2004] proposes an algorithm for query selection in situations where uncertainty over unidimensional utility functions is quantified probabilistically. At each stage, a myopically optimal query provides the largest reduction in the entropy of the joint distribution over utility values. Holloway and White [2003] consider sequentially optimal querying policies for a subclass of problems with additive utility functions and small sets of alternatives. The process is modeled as a special POMDP (see Section 4.3 below for a more detailed description).

While such methods strive to minimize the number of queries, they fail to consider the tradeoff between elicitation costs and improvement in decision quality. This is the topic of the next section.

3.3.3 Expected value of information

If uncertainty over utilities is quantified probabilistically, the value of a query can be computed by considering the values of updated belief states (one for each possible response), and weighting those values by the probability of corresponding responses. If a sequence of queries can be asked, finding the best elicitation policy is a sequential decision process, providing an optimal tradeoff between query costs (the burden of elicitation) and potentially better decisions due to additional information. However, such a policy is very difficult to compute; therefore, we first describe a myopic approach to choosing the next query.

Myopic EVOI Because of computational complexity of determining full value of a query, it is common to use myopic *expected value of information* (EVOI) to determine appropriate queries [Chajewska et al., 2000]. To reduce uncertainty about utility functions, the decision support system can ask questions about the user’s preferences. We will assume a finite set of available *queries* $Q = \{q_1, \dots, q_n\}$, and, for each query q_i — a set of possible user *responses* $R^i = \{r_1^i, \dots, r_m^i\}$. Responses to queries depend on the true user utility function u , but might be noisy. A general model that fits many realistic scenarios is a probabilistic *response model* $Pr(r_j^i|q_i, u)$, providing the probability of response r_j^i to the query q_i when the utility function is u . $Pr(r_j^i|q_i, \pi)$ will denote the probability of response r_j^i with respect to the density π over utility functions:

$$Pr(r_j^i|q_i, \pi) = \int_{u \in U} Pr(r_j^i|q_i, u)\pi(u)du. \quad (24)$$

Elicitation of preferences takes time, imposes cognitive burden on users, and might involve considerable computational and financial expense. Such factors can be modeled by assigning each query q_i a *query cost* c_i .⁶ In a Bayesian formulation of the elicitation process, expected gains in decision quality should outweigh elicitation costs.

Let’s recall that $EU(x, \pi)$ is the expected utility of outcome x when π is the probability distribution over utilities (see Eq. 23). Let $MEU(\pi)$ be the *maximum expected utility* of belief state π :

$$MEU(\pi) = \max_{x \in X} EU(x, \pi). \quad (25)$$

A response r to a query q provides information about the true utility function and changes our current beliefs π according to the Bayes’ rule:

$$\pi^r(u) = \pi(u|r) = \frac{Pr(r|u)\pi(u)}{Pr(r|\pi)}. \quad (26)$$

Thus, after response r , the maximum expected utility is $MEU(\pi^r)$. To calculate the value of a query, the *MEUs* of its possible responses should be weighed according to their likelihood. The *expected posterior utility* of the query q_i is:

$$EPU(q_i, \pi) = \sum_{r \in R^i} Pr(r|q_i, \pi) MEU(\pi^r). \quad (27)$$

The *expected value of information* of the query q_i is its expected posterior utility minus its current maximum expected utility:

$$EVOI(q_i, \pi) = EPU(q_i, \pi) - MEU(\pi). \quad (28)$$

EVOI of the query q_i denotes the gain in expected value of the ultimate decision. A *myopically optimal* querying strategy would always select a query whose EVOI is greatest, after accounting for query costs. A

⁶More generally, costs could depend on the true utility function, or be associated with responses.

sequentially optimal strategy would consider the value of future queries when computing the EVOI of the current query. Even though some query might be very costly in short term, it might be able to direct the elicitation process to good regions (in terms of decision quality) of the utility space which might otherwise remain unexplored by the myopic EVOI strategy. The myopic EVOI approach is more popular in practice (used in [Chajewska and Koller, 2000, Braziunas and Boutilier, 2005]) because computational requirements of sequential EVOI are often prohibitive.

Sequential EVOI An obvious way to minimize the shortcomings of myopic querying strategies is to perform a multistage lookahead. Unfortunately, such multistage search would have to be computed online (during the execution of the policy), which might seriously limit its benefits.

Another approach is to compute a sequentially optimal policy offline. Boutilier [2002] introduces the concept of preference elicitation as a POMDP that takes into account the value of future questions when determining the value of the current question. As before, we assume a system that makes decisions on behalf of a user; such a system has a fixed set of choices (actions, recommendations) whose effects are generally known precisely or can be modeled stochastically. The system interacts with a user in a sequential way; at each step it either asks a question, or determines that it has enough information about a user's utility function to make a decision. As each query has associated costs, the model allows the system to construct an optimal interaction policy which takes into account the trade-off between interaction costs and the value of provided information. The approach is discussed in more detail in Section 4.6.2.

4 Research on preference elicitation

In this section, we survey several fields in decision analysis, consumer research, and AI, where preference elicitation plays a central role. Imprecisely specified multiattribute utility theory (ISMAUT) is one of the earlier attempts to consider decision making under partial preference information in classical decision analysis. Its extensions to engineering design and configuration problems have been influential in spurring recent interest in preference elicitation among artificial intelligence researchers. Conjoint analysis and analytical hierarchy process (AHP) methods that were developed largely in isolation in the fields of marketing research and decision analysis also attempt to solve preference elicitation issues of general interest. We finish by providing an overview of some recent work in preference elicitation in AI.

4.1 ISMAUT

“Screening of Multiattribute Alternatives” [Sarin, 1977]
“An interactive procedure for aiding multiattribute alternative selection” [White et al., 1983].
”A Model of Multiattribute Decisionmaking and Trade-off Weight Determination under Uncertainty” [White et al., 1984]
“Ranking With Partial Information: A Method and an Application” [Kirkwood and Sarin, 1985]
“Partial Information, Dominance, and Potential Optimality in Multiattribute Utility Theory” [Hazen, 1986]
“A penalty function approach to alternative pairwise comparisons in ISMAUT” [Anandalingam and White, 1993]

One of the earlier attempts to consider decision making under partial preference information is the work on *imprecisely specified multiattribute utility theory*, or ISMAUT, by White et al. [1983, 1984], Anandalingam and White [1993]. A similar framework was proposed before by Fishburn [1964] and Sarin [1977]. Related research by, e.g., Kirkwood and Sarin [1985], Hazen [1986], Weber [1987], deals with similar issues, even though it is not customarily called ISMAUT.

ISMAUT applies to situations in which the utility function can be written in a normalized additive form, i.e., as a sum of weighted local value functions for each attribute. The decision maker has to choose from a finite set of multiattribute alternatives. The goal of ISMAUT is to restrict the set of alternatives to those that are not dominated by any other alternative, based on the prior information on local value functions, weights, and comparisons between pairs of alternatives. If the reduced alternative set is too big for the decision maker to make a choice, we should assess local value functions or weights more accurately, reduce the set of nondominated alternatives, and continue the process as long as is necessary for optimal alternative selection. An obvious drawback of this scheme is the lack of an intelligent query selection strategy to drive the elicitation process. In the following section, we discuss the research that considers querying strategies in ISMAUT-like elicitation settings.

Let A be the set of size m of available multiattribute alternatives, whose generic element is \mathbf{x} . Each alternative \mathbf{x} is a point in an n -dimensional consequence space: $\mathbf{x} = (x_1, x_2, \dots, x_n)$. The preference relation over A can be expressed by an additive utility function:

$$u(\mathbf{x}) = \sum_{i=1}^n w_i v_i(x_i) = \mathbf{w} \cdot \mathbf{v}(\mathbf{x}),$$

where \mathbf{w} is a vector of weights, and $\mathbf{v}(\mathbf{x})$ is a vector of local value function values of alternative \mathbf{x} .

The model can incorporate three types of prior information (or responses to utility queries): comparison of attribute weights w_i , information about local value functions $v_i(\cdot)$, expressed by sets of linear inequalities, and pairwise preference statements about alternatives in the set A . In particular,

- 1 Knowledge about relative importance of the tradeoff weights (“Color is more important than screen size”) or bounds on their values (“This attribute’s weight is between 0.5 and 1”) ⁷ allows us to define a feasible subset $W \subseteq \{\mathbf{w} \in \mathbb{R}^n : w_i \geq 0, \sum_i w_i = 1\}$ of all possible weights via linear constraints.
- 2 Similar to statements about weights, ISMAUT allows us to model information about individual local value functions by means of linear constraints. If the third attribute is computer’s speed, and the user prefers faster computers, ceteris paribus, then $v_3(\text{‘fast’}) \geq v_3(\text{‘slow’})$. The user might also be able to provide bounds for local values of specific attribute levels (e.g., $v_3(\text{‘fast’}) \in [0.3; 0.7]$). Such linear constraints define the spaces V_1, \dots, V_n of possible local value functions.
- 3 Finally, even if the user is unable to select the best alternative right away, she might be able to compare some pairs of alternatives. Let J be the set of such comparisons: for each $j \in J$ there is a pair $(\hat{\mathbf{x}}^j, \mathbf{x}^j) \in A \times A$ if and only if the user has specified that $\hat{\mathbf{x}}^j \succeq \mathbf{x}^j$. The set of comparisons J can be used to impose further restrictions on the weight space, because $j \in J$ implies $\mathbf{w} \cdot \mathbf{v}(\hat{\mathbf{x}}^j) \geq \mathbf{w} \cdot \mathbf{v}(\mathbf{x}^j)$.

All this prior information defines the set C of feasible weights and local value functions. More precisely, the tuple $\langle \mathbf{w}, v_1, \dots, v_n \rangle \in C$ if and only if

$$\begin{aligned} \mathbf{w} &\in W, \\ v_i(\cdot) &\in V_i, \text{ for all } i = 1, \dots, n, \\ \mathbf{w} \cdot [\mathbf{v}(\hat{\mathbf{x}}^j) - \mathbf{v}(\mathbf{x}^j)] &\geq 0, \text{ for all } j \in J. \end{aligned}$$

⁷Although many authors talk about the “importance” of attributes, we should be aware that weights are nothing more than scaling factors. The statements about weights are nonetheless meaningful: $w_i \geq w_j$ means that outcome $(x_i^\top, \mathbf{x}_{iC}^\perp)$ is preferred to $(x_j^\top, \mathbf{x}_{jC}^\perp)$, and $w_i \in [0.5; 1]$ means that $u(x_i^\top, \mathbf{x}_{iC}^\perp) \in [0.5; 1]$.

The additive structure of utility functions allows us to use the set C to eliminate dominated alternatives. First, we define the binary relation $R(C) \subseteq A \times A$ as follows:

$$(\hat{\mathbf{x}}, \mathbf{x}) \in R(C) \iff \mathbf{w} \cdot [\mathbf{v}(\hat{\mathbf{x}}) - \mathbf{v}(\mathbf{x})] \geq 0, \text{ for all } \langle \mathbf{w}, v_1, \dots, v_n \rangle \in C.$$

This means that $\hat{\mathbf{x}} \succeq \mathbf{x}$ if and only if $(\hat{\mathbf{x}}, \mathbf{x}) \in R(C)$.⁸ When $\mathbf{w} \cdot \mathbf{v}(\mathbf{x}') = \mathbf{w} \cdot \mathbf{v}(\mathbf{x})$ for all $\langle \mathbf{w}, v_1, \dots, v_n \rangle \in C$, then \mathbf{x} and \mathbf{x}' are said to be equal (with respect to C).

The set of *nondominated* alternatives $ND(C)$ can be computed using the relation $R(C)$: $\hat{\mathbf{x}}$ is non-dominated if there is no nonequal alternative \mathbf{x} such that $(\mathbf{x}, \hat{\mathbf{x}}) \in R(C)$. We should note that without prior information, C contains all possible weights and value functions, and $ND(C)$ is equal to the Pareto-optimal set of alternatives. $ND(C)$ is important because the most preferred alternative has to be in it. The goal of ISMAUT is to reduce the set of nondominated alternatives until the user can select the optimal one. More information about the possible local value functions and weights reduces the set C , increases the binary relation $R(C)$, and reduces the size of $ND(C)$:

$$C \subseteq C' \implies R(C') \subseteq R(C) \wedge ND(C) \subseteq ND(C').$$

The set $ND(C)$ can be computed from $R(C)$ in polynomial time in size of the alternative set A , because for each alternative, we need to check that no other alternative is preferred. The central computational task is therefore to compute $R(C)$ from C . Recall that $(\hat{\mathbf{x}}, \mathbf{x}) \in R(C)$ if and only if $\mathbf{w} \cdot [\mathbf{v}(\hat{\mathbf{x}}) - \mathbf{v}(\mathbf{x})] \geq 0$, for all $\langle \mathbf{w}, v_1, \dots, v_n \rangle \in C$. This amounts to verifying that

$$\min_{\langle \mathbf{w}, v_1, \dots, v_n \rangle \in C} \mathbf{w} \cdot [\mathbf{v}(\hat{\mathbf{x}}) - \mathbf{v}(\mathbf{x})] \geq 0. \quad (29)$$

If $J = \emptyset$, then there is no prior information about pairwise preferences between alternatives; the constraint set C contains all the weights in W , and value functions in V_1, \dots, V_n . In this case, researched in White et al. [1983], Eq. 29 becomes

$$\min_{\mathbf{w} \in W} \left(\sum_{i=1}^n w_i \min_{v_i \in V_i} [v_i(\hat{x}_i) - v_i(x_i)] \right) \geq 0. \quad (30)$$

This problem can be solved in a straightforward manner by $n + 1$ linear programs.

Adding constraints to the set J complicates the solution, because the weight and value function constraints get tied together by comparisons of global alternatives. Anandalingam and White [1993] propose a general penalty function method for determining membership in $R(C)$. However, since such methods often suffer from slow convergence and ill-conditioning (see, e.g., [Bazaraa and Shetty, 1979]), two special cases of approximating $R(C)$ were analyzed by White et al. [1984], and Anandalingam and White [1993].

Let C^{min} be a set of tuples $\langle \mathbf{w}, v_1, \dots, v_n \rangle$ such that $\mathbf{w} \in W$, $v_i \in V_i$, and for all $j \in J$,

$$\sum_{i=1}^n w_i \min_{v_i \in V_i} [v_i(\hat{x}_i^j) - v_i(x_i^j)] \geq 0. \quad (31)$$

Similarly, let C^{max} be a set of tuples $\langle \mathbf{w}, v_1, \dots, v_n \rangle$ such that $\mathbf{w} \in W$, $v_i \in V_i$, and for all $j \in J$,

$$\sum_{i=1}^n w_i \max_{v_i \in V_i} [v_i(\hat{x}_i^j) - v_i(x_i^j)] \geq 0. \quad (32)$$

⁸In this case, the relation \succeq denotes derived, or provable preference, rather than “true” preference.

Inequalities 31 and 32 impose linear constraints on weights. Therefore, solving for $R(C^{min})$ and $R(C^{max})$ is straightforward, since this case is similar to the $J = \emptyset$ case described above.

How can C^{min} and C^{max} be used to approximate $R(C)$? From Eq. 31 and 32, we note that $C^{min} \subseteq C \subseteq C^{max}$, and therefore, $R(C^{max}) \subseteq R(C) \subseteq R(C^{min})$. Without computing $R(C)$, we can infer two facts:

$$\begin{aligned}(\hat{\mathbf{x}}, \mathbf{x}) \in R(C^{max}) &\implies (\hat{\mathbf{x}}, \mathbf{x}) \in R(C), \\(\hat{\mathbf{x}}, \mathbf{x}) \notin R(C^{min}) &\implies (\hat{\mathbf{x}}, \mathbf{x}) \notin R(C).\end{aligned}$$

Therefore, it is only when $(\hat{\mathbf{x}}, \mathbf{x}) \notin R(C^{max})$ or $(\hat{\mathbf{x}}, \mathbf{x}) \in R(C^{min})$ that membership in $R(C)$ needs to be determined directly.

Hazen [1986] and Weber [1987] point out that the set of nondominated alternatives is not the same as the set of *potentially optimal* alternatives, which is a subset of $ND(C)$. It is possible that an alternative is not dominated by any other alternative, but is dominated by a collection of alternatives (Weber [1987] calls this *mixed dominance*). Or, alternatively, a potentially optimal alternative always has a feasible *witness* utility function for which it is an optimal alternative. When local value functions are known, the witness weight vector can be found by solving a linear program.

4.2 Engineering design and configuration problems

"Multiobjective Intelligent Computer-Aided Design" [Sykes and White, 1991]

"Preference-Directed Design" [D'Ambrosio and Birmingham, 1995]

One field in which applications and extensions of ISMAUT have been proposed is *engineering design*. Design is a multidisciplinary area with no precise definition. Generally, any problem of designing a complex system that has to comply to some performance requirements and satisfy operational constraints can be regarded as an engineering design problem. Examples include communication networks, computer systems, bridges, etc. Some areas within engineering design that have tight connections to AI are *AI in design* (AID), *knowledge-based design systems* (KBDS), and *intelligent computer-aided design* (ICAD) [Brown and Birmingham, 1997]. As we shall see, *configuration design* [Wielinga and Schreiber, 1997] is a particularly relevant formalization of the problem with regard to preference elicitation. A recent report by the Board on Manufacturing and Engineering Design [2001] stresses the importance of the decision-theoretic approach in engineering design.

The paper of Sykes and White [1991] on *multiobjective intelligent computer-aided design* (MICAD) extends the ideas of ISMAUT to the problem of configuration design. The design process is viewed as a combination of the progressively acquired *preferential component* and an *a priori operational component*. MICAD thus combines iterative capture of user preferences with the search in the constrained space of feasible designs. Preference elicitation can be directed toward promising (and feasible) regions of the design space, thus avoiding the cost of wasted elicitation effort. On the other hand, the search for optimal designs can be substantially facilitated by preference information.

Intuitively, a configuration problem is that of "configuring" a system. A decision to be taken consists of a number of components, or aspects, which interact in complex ways to produce an outcome. A typical example is configuring a computer system from a set of components — choosing a processor, compatible memory, peripherals, etc. Possible *configurations* are restricted by hard feasibility constraints. The optimal configuration depends on user preferences; however, those preferences are expressed over *features* or *attributes* of configurations (or *designs*). For example, a user might want a "reliable home computer", which

is a point in the *feature space* (or *performance space*), rather than *configuration space*. A mapping from configuration space to feature space induces indirect preferences over configurations.

Let $\mathbf{T} = T_1 \times \dots \times T_m$ be a multiattribute configuration space, where each T_i is a set of components to choose from. As before, the outcome, or feature, space will be denoted as \mathbf{X} . Components represent controllable aspects of a design problem, whereas configuration features allow for a direct expression of user preferences. Feasible configurations $\mathbf{T}^F \subseteq \mathbf{T}$ form a subset of the configuration space; they could be specified by a set of rules, logical formulas, or using *constraint satisfaction problem* (CSP) formulations.⁹ A *performance function*¹⁰ $f : \mathbf{T} \mapsto \mathbf{X}$ provides a mapping from configuration space to feature space. The problem is complicated by the fact that this function $f(\cdot)$ might not have any useful mathematical properties (such as continuity, monotonicity or invertibility), and might not be expressed in closed form. Determination of $f(t)$, $t \in \mathbf{T}$ might also be costly and require expert analysis or simulation.

The problem of engineering design can thus be summarized follows: given a set of components and features, a set of operational constraints on configurations, a performance function, and a preference relation over the outcome space, find an optimal feasible configuration. This general problem is addressed in [Sykes and White, 1991, D’Ambrosio and Birmingham, 1995] using ISMAUT, and in [Boutilier et al., 1997] using CP-nets.

Sykes and White [1991] investigate direct application of ISMAUT ideas to the design process. It is assumed that local value functions are known, so only weights are uncertain. Information about weights can be queried from the user in two ways, already described above: (1) the user can provide direct information about weights, expressed as linear constraints; (2) the user can compare pairs of designs to induce linear inequalities in the weight space (this requires solving Eq. 30). Preference elicitation can occur at any time during an iterative design process. MICAD is presented as a general framework for interactive preference elicitation and search in the space of designs. It is assumed that the search proceeds in stages, at which a finite set of designs is available for the user to evaluate. Two crucial issues are not directly addressed: how to select a set of designs (from potentially exponential number of possibilities) at each stage, and what query selection strategy to follow when eliciting user preferences.

D’Ambrosio and Birmingham [1995] tackle the first issue. They formulate the design engineering problem as a constrained optimization problem. The objective function is an incomplete value function created by pairwise ranking a random sample of design alternatives. Operational constraints are modeled as a CSP. Therefore, CSP solution techniques, such as constraint network decomposition and constraint propagation, can be harnessed to facilitate a branch-and-bound search for optimal designs.

4.3 Extensions of ISMAUT

“Q-Eval: Evaluating Multiple Attribute Items using Queries” [Iyengar et al., 2001]
 “Polyhedral Sampling for Multiattribute Preference Elicitation” [Ghosh and Kalagnanam, 2003]
 “Question Selection for Multiattribute Decision-aiding” [Holloway and White, 2003]

Classical ISMAUT is mostly concerned with narrowing the set of alternatives to a manageable size using partial preference information. The following papers address an important issue of how to select queries in a sequential elicitation process. Like in ISMAUT, the additive utility function over attributes is assumed.

The Q-Eval algorithm of Iyengar et al. [2001] asks the user to compare pairs of selected alternatives, and uses the responses to refine the preference model. The local value functions are specified precisely,

⁹A constraint satisfaction problem deals with finding a feasible assignment to a set of variables subject to a set of constraints. Dechter [2003] provides a detailed overview of CSP algorithms and models.

¹⁰In [Boutilier et al., 1997], the performance function is called a *causal model*, because it is expressed by a set of logical rules.

so utility function uncertainty is represented by linear constraints on the weight space. Given a set of alternatives, the authors address the issue of which pair to present to the user for ranking in a sequential elicitation process. Each response to the query “Is \hat{x} preferred to x ?” adds a linear constraint which reduces the region of feasible weights W . Since a response is not known beforehand, the authors advocate a heuristic of choosing the query that would come closest to bisecting the space of feasible weights. The rationale for this querying strategy is to shrink the space of possible weights as quickly as possible.

The implementation of Q-Eval employs a number of approximations to ensure practical online performance. First, the number of alternative pairs considered is pruned based on the normal distance of corresponding hyperplanes to the “center” of the region W . Intuitively, hyperplanes close to the center are good candidates for bisecting the region W equally. The notion of center used throughout the paper is that of *prime analytic center*, which is the point that maximizes the sum of log distances to the irredundant hyperplanes defining the region. In case a decision has to be made with uncertain information, the center serves as a representative weight vector. Queries that were not pruned in the previous step are then evaluated based on the volumes of the resulting polytopes (the best query leads to the most equal partition of the weight space). The volumes are approximated by the size of the tightest axis-orthogonal bounding rectangle.

Ghosh and Kalagnanam [2003] consider the same problem and propose to use sampling for determining the center of the weight region W . In particular, they use a hit-and-run sampling technique that employs a Markovian random walk defined on the set W with a uniform stationary distribution. The advocated querying strategy is to ask a query whose corresponding hyperplane is orthogonal to the longest line segment contained in W . The method is quite *ad hoc*, but works fast in practice.

The two query selection methods described above try to minimize the number of queries by shrinking the region of possible weights as fast as possible. However, they do so myopically, without considering the value of a *sequence* of queries. Holloway and White [2003] present a POMDP model for sequentially optimal elicitation in an ISMAUT-like setting, where uncertainty about utility functions is specified by linear constraints on weight vectors.

The state space in this POMDP model is the (uncountable) collection of all subsets of tradeoff weights $\{\mathbf{w} \geq 0 : \sum w_i = 1\}$. Intuitively, a system is in state W , if W is the largest region of the weight space constrained by previous elicitation responses. Actions are binary queries asking to compare pairs of alternatives, and observations are *yes/no* answers to such queries. It is assumed that there is no noise in user responses. The observation function is the probability of getting a response r to the query q when the true tradeoff weight vector lies in the set W ; a uniform probability distribution over the weight space is assumed, although more general probability models could be accommodated. The process moves from one state to another as the feasible weight region shrinks due to linear constraints imposed by responses to queries.

To define the cost structure, we need to introduce a notion of the solution partition, which divides the set of all weights into convex regions where one alternative dominates all others. Formally, if A is the set of alternatives, a *solution partition* is $\{W_a : a \in A\}$, where

$$W_a = \{\mathbf{w} : \mathbf{w} \cdot \mathbf{v}(\mathbf{x}_a) \geq \mathbf{w} \cdot \mathbf{v}(\mathbf{x}_{a'}) \forall a' \in A\}. \quad (33)$$

Elicitation goals can be encoded using the cost model. Let T be a maximum number of queries that can be asked (thus, we consider a finite-horizon POMDP). For a given weight vector set W , $c(W, q)$ is the cost of asking the query q , and $\bar{c}(W)$ is the terminal cost of ending up with the set W after all questions have been asked. If the goal is to ask as few queries as possible to determine an optimal alternative, then we can set $\bar{c}(W) = 0$, $c(W, q) = 0$ if there is $a \in A$ such that $W \subseteq W_a$, and $c(W, q) = 1$ otherwise. The optimal policy of this POMDP will ask the fewest queries possible until it finds the smallest $t \leq T$ such

that $W^{(t)} \subseteq W_a$ for some $a \in A$. One can similarly define a simple cost function for minimizing the expected uncertainty of knowing the problem’s solution after T questions.

Since the POMDP described above is hard to solve in the most general form (although ideas from [Boutilier, 2002] would be applicable here as well), the authors concentrate on the case in which restrictions on the cost function guarantee a finite, piecewise linear, representation of the POMDP value function. This is possible if the cost function depends on W only through a finite probability distribution $p^W(\cdot)$ over alternatives. For each $a \in A$, $p^W(a)$ is the probability that alternative a is optimal, given that the true weight vector is in W . Such cost functions are too restrictive for POMDPs that model the optimal tradeoff between elicitation costs and expected improvement in decision quality; however, they can be used to achieve the two goals mentioned in the previous paragraph.¹¹

Holloway and White [2003] do not perform empirical validation of the approach or provide a suitable POMDP solution algorithm. Quite unrealistically, the authors also assume perfect responses to queries. Nevertheless, it is the first attempt to describe a model for sequentially optimal query selection in ISMAUT problems.

4.4 Conjoint analysis

“Conjoint Measurement for Quantifying Judgmental Data” [Green and Rao, 1971]
 “Conjoint Analysis in Consumer Research: Issues and Outlook” [Green and Srinivasan, 1978]
 “Conjoint Analysis in Marketing: New Developments with Implications for Research and Practice” [Green and Srinivasan, 1990]
 “Polyhedral Methods for Adaptive Choice-Based Conjoint Analysis” [Toubia et al., 2004]
 “Fast Polyhedral Adaptive Conjoint Estimation” [Toubia et al., 2003]

Since the original paper by Green and Rao [1971], conjoint analysis has become a major area in marketing research.¹² Conjoint analysis is a set of techniques for measuring consumer tradeoffs among multi-attribute products and services. Despite differences in terminology and methodology, conjoint analysis and multiattribute decision analysis (in particular, ISMAUT) deal with similar issues in preference elicitation and modeling.

The goal of conjoint analysis is to decompose consumer preferences over multiattribute *products* (or *profiles*) into component preferences over attributes in order to predict *aggregate* consumer behavior, explain preferences for current products, visualize market segmentation, and help design new products. Thus, the emphasis is generally on predictive and descriptive, rather than prescriptive aspects of consumer behavior.

Usually, an additive utility function is assumed — the total value of a product is the sum of partial contributions (*partworths*) of individual attributes (*features*). Formally, let $y^j = u(\mathbf{x}^j)$ be a specified rating of the product \mathbf{x}^j . A general conjoint analysis model is

$$y^j = \sum_{i=1}^t v_i z_i^j, \tag{34}$$

where z_i^j are input variables, y^j is a dependent output variable, and v_i are parameters to be estimated. Input variables z_i^j depend on the attributes of the product \mathbf{x}^j :

¹¹For example, the problem of minimizing the number of queries can be encoded by setting $\bar{c}(p^W) = 0$, $c(p^W, q) = 0$ if there exists $a \in A$ such that $p^W(a) = 1$, and $c(p^W, q) = 1$ otherwise.

¹²[Green and Srinivasan, 1978] and [Green and Srinivasan, 1990] are key historical surveys of conjoint analysis.

- For continuous attributes whose value is monotonically increasing, $z_i = x_i$. If all attributes are like that, then the model reduces to the familiar linear value function $u(\mathbf{x}) = \sum_i w_i x_i$, and parameters v_i can be viewed as weights w_i .
- For continuous attributes whose local value functions are substantially nonlinear, several z_i variables can be used for approximation. In a case of quadratic function for attribute i , two z variables are introduced: one equal to x_i , and the other equal to x_i^2 . Such local value models are quite common in conjoint analysis (one example is the *ideal-point model*, where local preference increases quadratically until some ideal-point level, and decreases after that).
- For discrete binary attributes with two levels x_i^\top and x_i^\perp , we set $z_i = 1$ if $X_i = x_i^\top$, and $z_i = 0$ if $X_i = x_i^\perp$. Then an estimated parameter v_i can be thought of as a local value of the best level of x_i .
- Discrete attributes with k levels are converted into $k - 1$ binary “dummy” attributes. Constraints on indicator variables z_i are added to ensure consistency of representation.

Given preference information about whole products (such as ordinal or cardinal product ranking, comparison, or preferred choice from a set of products), some form of regression is used to find parameters that are *most consistent* with specified preferences, which are usually aggregate. For example, a common type of application is to elicit preferences over full profiles using a rating or ranking scale, and then estimate attribute partworths by least-squares regression. The underlying assumption is that ranking or rating full products is easier than providing attribute partworths, as long as the number of attributes is small.

Many aspects of preference elicitation considered in this report have their equivalents in conjoint analysis, too. Approaches are differentiated according to data collection formats (i.e., “query types”), question design (“query selection”), and parameter estimation procedures (“decision making with incomplete information”). The most common data collection format is full profile evaluation, where a user is asked to order all products (*stimuli*) in a given set, or provide a metric rating of each stimulus. Of course, the user’s burden grows dramatically with the size of stimulus set. Some methods therefore employ partial profile evaluations. *Choice-based conjoint analysis* (CBC) is a popular compromise technique, where instead of ranking all profiles, a user is asked to choose the most preferred from a given a set. *Metric paired-comparison* format asks to consider only pairs of profiles, but expects quantitative answers regarding relative preference.¹³

Until recently, most applications of conjoint analysis either presented the same questions to all respondents, blocked them across sets of respondents, chose randomly, or adapted them based on responses from prior respondents. Adaptive question design for individual respondents in the manner of ISMAUT was first considered by Toubia et al. [2003, 2004] in metric paired-comparison and CBC settings. This new approach, termed the *polyhedral method*, works by iteratively constraining the polyhedron of feasible subutility (partworth) values. The attributes are discrete and binary (multilevel attributes can be represented using dummy variables), so each product is represented by a point in the space of attribute partworths. In CBC, binary comparison questions result in a separating hyperplane that cuts the polyhedron of feasible subutilities. More generally, a respondent is presented with a set of products, and asked to choose one of them. A choice set of size k defines $k(k - 1)/2$ possible hyperplanes; for each of k choices available, $k - 1$ hyperplanes determine the new polyhedron.

In polyhedral methods, the goal is to reduce the size of uncertainty polyhedron as fast as possible. Questions are designed to partition the polyhedron into approximately equal parts; in addition, shape heuristics are used to favor cuts that are perpendicular to long axes. Since the problem is computationally hard, many approximations similar to Q-Eval [Iyengar et al., 2001] are employed. The polyhedron’s volume is approximated by a bounding ellipsoid, and its center by the analytic center. Then, k points at which $k/2$ longest

¹³In practice, a user is usually provided with a set of qualitative choices specifying by how much product x is preferable to product y (e.g., “I like x much more than y ”, “I like x a little more than y ”, “I like x as much as y ”, etc.); these choices are then converted to a quantitative scale.

axes intersect the polyhedron are used to select k profiles for the next choice-based query. This technique is extended to metric paired-comparison queries in Toubia et al. [2003].

Conjoint analysis and decision analysis have largely developed in parallel, without much interaction. However, recent emphasis on sequential preference elicitation in both fields presents opportunities for fertile interaction. Conjoint analysis offers a variety of query formats that have been validated in practice, and many experimental domains in consumer research. Its limitations include reliance on full profile queries,¹⁴ which work only for products with a few (usually less than ten) attributes, common assumptions of attribute independence, and lack of integration of preference elicitation and product feasibility constraints.

4.5 Analytic hierarchy process

“A scaling method for priorities in hierarchical structures” [Saaty, 1977]
 “The analytic hierarchy process” [Saaty, 1980]
 ”Preference Ratios in Multiattribute Evaluation (PRIME)–Elicitation and Decision Procedures under Incomplete Information” [Salo and Hämäläinen, 2001]

Analytic hierarchy process (AHP) is an alternative method of decision analysis developed by Saaty [1977, 1980]. The main ideas of the AHP method can be explained in comparison to additive value theory (see [French, 1986]), although the connection between the two approaches was developed well after the original work on AHP.

The problem is to select the best alternative from the set of m multiattribute alternatives $\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^m$ under certainty. Each alternative is measured against n attributes: $\mathbf{x}^i = (x_1^i, x_2^i, \dots, x_n^i)$. The value function is represented as a weighted sum of strictly positive local value functions $v_i(\cdot)$:

$$v(\mathbf{x}^i) = \sum_k^n w_k v_k(x_k^i) = \sum_k^n w_i v_k^i, \quad (35)$$

where v_k^i is the local value of the i th alternative on the k th attribute. The weights and local value functions are not normalized to $[0; 1]$.

The main difference between AHP and classical decision analysis lies in the elicitation of weights and local value functions. Instead of direct responses regarding attribute weights and local value functions, AHP assumes that a user can instead provide all the entries of the so-called *positive reciprocal matrices*. For each attribute k , the entries can be thought of as ratios of local value functions:

$$R_k = \begin{pmatrix} 1 & r_k^{12} & \dots & r_k^{1m} \\ 1/r_k^{12} & 1 & \dots & r_k^{2m} \\ \vdots & \vdots & \dots & \vdots \\ 1/r_k^{1m} & 1/r_k^{2m} & \dots & 1 \end{pmatrix} = \begin{pmatrix} v_k^1/v_k^1 & v_k^1/v_k^2 & \dots & v_k^1/v_k^m \\ v_k^2/v_k^1 & v_k^2/v_k^2 & \dots & v_k^2/v_k^m \\ \vdots & \vdots & \dots & \vdots \\ v_k^m/v_k^1 & v_k^m/v_k^2 & \dots & v_k^m/v_k^m \end{pmatrix}, \quad (36)$$

where r_k^{ij} is the *ratio* of local values v_k^i and v_k^j . Besides k attribute matrices, an additional matrix R is elicited to provide information about relative importance of attributes.

$$R = \begin{pmatrix} w_1/w_1 & w_1/w_2 & \dots & w_1/w_n \\ w_2/w_1 & w_2/w_2 & \dots & w_2/w_n \\ \vdots & \vdots & \dots & \vdots \\ w_n/w_1 & w_n/w_2 & \dots & w_n/w_n \end{pmatrix}. \quad (37)$$

¹⁴There are methods of conjoint analysis that do not employ full product comparisons, but they are less popular and not as well grounded theoretically from the decision theory perspective.

The entries of R can be interpreted as ratios of attribute weights.

Given the entries of the positive reciprocal matrices, AHP derives the weights and local value functions for the attributes. If the matrix entries were consistent, such derivation would amount to solving a simple system of linear equations. In a likely case of inconsistent entries, the local value functions and weights are estimated using one of several averaging techniques (eigenvector-based estimation is commonly advocated). The alternatives are ultimately ranked by the resulting additive value function.

The key issue is eliciting the positive reciprocal matrices. For attribute matrices, the user is asked to compare pairs of alternatives on each attribute.¹⁵ For a pair of alternatives \mathbf{x}^i and \mathbf{x}^j compared on attribute k , the ratio r_k^{ij} is 1, if \mathbf{x}^i is *equally* preferred to \mathbf{x}^j , 3 — *weakly* preferred, 5 — *strongly* preferred, 7 — *demonstrably* preferred, and 9 — *absolutely* preferred. The weight matrix is elicited by a similar process — one attribute can be “equally important”, “weakly more important”, “strongly more important”, “demonstrably more important”, and “absolutely more important” than another.

Issues such as elicitation costs, decision making with incomplete information, and query selection criteria are as important in the AHP as in classical decision theory. Salo and Hämäläinen [1995, 2001, 2004] maintain that decisions should be made with incomplete information if elicitation costs outweigh potential improvement in decision quality. Preference uncertainty is described by bounds on value function ratios (i.e., the entries of positive reciprocal matrices). Several decision criteria are discussed, and “central values” approach (see Section 3.2.1) favored on the grounds of empirical simulations. Query selection is not addressed.

AHP is a controversial method (see, e.g., [French, 1986, Salo and Hämäläinen, 1997] for some criticisms of AHP). While quite popular in practice,¹⁶ it is not as well grounded theoretically as classical decision theory. One problem is that while local value functions are interval scales, the construction of positive reciprocal matrices assumes that they are *ratio* scales;¹⁷ AHP fails to provide an axiomatic basis for such a strong assumption. Elicitation of matrix entries is also problematic, since it hard to provide an exact semantic meaning to the AHP queries. The nine-point scale is a source of further controversy. If level 1 of an attribute is absolutely preferred to level 2, and level 2 is absolutely preferred to level 3, then the ratio of level 1 and level 3 should be $9 \times 9 = 81$. However, the scale allows only numbers from 1 to 9. Finally, AHP violates the principle of independence of irrelevant alternatives (i.e., the principle that ranking between two alternatives should be independent of other available alternatives).

4.6 Preference elicitation in AI

4.6.1 Minimax regret approach

”Incremental Utility Elicitation with the Minimax Regret Decision Criterion” [Wang and Boutilier, 2003]

”Cooperative Negotiation in Autonomic Systems using Incremental Utility Elicitation” [Boutilier et al., 2003a]

”New Approaches to Optimization and Utility Elicitation in Autonomic Computing” [Patrascu et al., 2005]

”Constraint-based Optimization with the Minimax Decision Criterion” [Boutilier et al., 2003b]

”Regret-based Utility Elicitation in Constraint-based Decision Problems” [Boutilier et al., 2005]

¹⁵*Ceteris paribus* with respect to remaining attribute values should certainly be assumed, although such issues, and many other, are often skirted in AHP literature.

¹⁶Among the main reasons for AHP popularity is the relative simplicity of elicitation queries: to obtain a total ranking of multi-attribute alternatives, a user is only asked to provide a qualitative comparison between *pairs* of attributes.

¹⁷Functions on a ratio scale are unique up to positive scaling.

Minimax regret criterion can be used both for making robust decisions under strict uncertainty and for driving an elicitation process. Contrary to the methods in the previous section, the quality (difference from optimal) of a minimax regret optimal decision can be bounded; these bounds can be tightened with further elicitation effort. Minimax regret methods have been applied to several areas of AI, including auctions [Wang and Boutilier, 2003], autonomic computing [Boutilier et al., 2003a, Patrascu et al., 2005], combinatorial auctions [Boutilier et al., 2004b], and constrained configuration problems [Boutilier et al., 2003b, 2005].

Wang and Boutilier [2003] consider a simple problem with a flat outcome space and binary standard gamble queries. A response to a query results in a new decision situation with a new level of minimax regret. The (myopic) value of a query is a function of response values. The authors consider three ways of combining response values: maximin improvement (select the query with the best worst-case response), average improvement (select the query with the maximum average improvement), and expected improvement (select the best query based on improvements weighted by the likelihood of responses). It turns out that the expected improvement criterion, combining a Bayesian query selection strategy and a robust minimax regret decision criterion, performs best experimentally and is not subject to stalling — the situation when no query improves the minimax regret level. Using binary standard gamble queries, the querying strategy can be optimized analytically.

Boutilier et al. [2003b] address the problem of choosing the best configuration from the set of feasible configurations encoded by hard constraints. It is assumed that preferences over configurations can be represented by a GAI utility function; however, this function is imprecisely specified by bounds on GAI subutility function values. The authors propose the use of minimax regret as a suitable decision criterion and investigate several algorithms based on mixed integer linear programming to compute regret-optimizing solutions efficiently.

In [Boutilier et al., 2005], the authors concentrate on the utility elicitation aspect and provide an empirical comparison of minimax regret reduction strategies in GAI utility models, where uncertainty over utilities is expressed by bounds on local factor values. The objective is to refine utility uncertainty and reduce minimax regret with as few queries as possible. The queries are *bound* queries: the user is asked whether a specific local utility parameter lies above a certain value. A positive response raises the lower bound, while a negative response lowers the upper bound of a local subutility value.

The *halve largest gap* (HLG) elicitation strategy recommends a query at the midpoint of the bound interval of the GAI factor setting with the largest gap between upper and lower bounds. HLG uniformly reduces uncertainty over the entire utility space and therefore provides the best theoretical minimax regret reduction guarantees. It is related to polyhedral methods (with rectangular polytopes) in conjoint analysis which attempt to maximally reduce uncertainty with each query. Another, *current solution* (CS), strategy, uses heuristics to focus on *relevant* parts of the utility space and works better in practice. CS relies on two special outcomes that are directly involved in calculating the regret level: \mathbf{x}^* , the minimax optimal configuration, and \mathbf{x}^w , the *witness* point that maximizes the regret of \mathbf{x}^* . The CS strategy considers only local factor settings that are part of these two special outcomes and asks about the one with the largest gap. A few other heuristic strategies are also tested in experiments.

The minimax regret criterion can also be applied to a completely different domain of autonomic computing [Boutilier et al., 2003a, Patrascu et al., 2005]. To solve the problem of optimal resource allocation one needs to know the utility of different levels of resource applied to the distributed computing elements. Unfortunately, even a single evaluation of the utility function is very costly. Patrascu et al. [2005] investigate how to sample a monotonic non-decreasing utility function with a continuous unidimensional domain using strategies similar to CS and HLG.

4.6.2 Bayesian approach

"Utilities as Random Variables: Density Estimation and Structure Discovery" [Chajewska and Koller, 2000]
"Making Rational Decisions Using Adaptive Utility Elicitation" [Chajewska et al., 2000]
"A POMDP Formulation of Preference Elicitation Problems" [Boutilier, 2002]
"Local Utility Elicitation in GAI Models" [Braziunas and Boutilier, 2005]

If uncertainty about utility functions can be quantified probabilistically, then one can design preference elicitation strategies that optimally balance the tradeoff between elicitation effort and the impact of information on the decision quality. Until recently, this approach has been explored very little. In this section, we take a look at some of the attempts to solve this problem by AI researchers.

Myopic EVOI The work of [Chajewska et al., 2000] was arguably the first to adopt a consistent Bayesian view of the preference elicitation problem. If the utility function is not fully known, it is treated as a random variable drawn from the prior distribution [Chajewska and Koller, 2000]. The value of a decision in an uncertain situation is computed by taking an expectation over all possible utility functions. Furthermore, the value of a query is simply its expected value of information.

The proposed framework leads to a simple elicitation algorithm. At each step, the query with the highest EVOI is asked, and the distribution over utilities is updated based on user responses. The process stops when the expected value of a decision meets some termination criteria. Because the sequential EVOI (which takes into consideration all possible future questions and answers) is hard to compute, the value of a query is approximated by the *myopic* EVOI (see Eq. 28).

In the prenatal diagnosis decision model described in the paper, the outcome space of size n is discrete and unstructured (flat). Therefore, the space of all utility functions can be represented by an n -dimensional unit hypercube. A multivariate Gaussian distribution (restricted to $[0;1]$) is used to model the prior over utilities. After a binary standard gamble query ("Is utility of outcome x greater than p ?"), the resulting posterior becomes a truncated Gaussian, which is then *approximated* by a new multivariate Gaussian distribution. Experimental results on the domain with 108 outcomes show that very few queries are needed to reduce the expected utility loss below a small threshold.

Braziunas and Boutilier [2005] also adopt a myopic approach to choosing the next query in eliciting parameters of GAI models. In this case, EVOI computation is facilitated by the additive structure of GAI utilities. The uncertainty over utilities is quantified via independent priors over local value function parameters. In such a case, an appropriate form of query is "Is local utility of suboutcome x_i greater than l ?". Such queries are *local* queries, because they ask a user to focus on preferences over a (usually small) subset of attributes; the values of remaining attributes do not have to be considered. The authors show that the best myopic query can be computed analytically if the prior information over local utility parameters is specified as a mixture of uniform distributions [Boutilier, 2002]. Such mixtures fit nicely with the type of queries that result in axis-parallel density "slices", because the posterior distribution after a response to a query remains a mixture of uniforms. It is therefore possible to maintain an exact density over utility parameters throughout the elicitation process.

Preference elicitation as a POMDP To overcome the shortcomings of myopic EVOI approaches, the preference elicitation problem can be modeled as a POMDP [Boutilier, 2002]. The state space of the preference elicitation POMDP is the set of possible utility functions U ; actions can be either queries about a user's utility function Q or terminal decisions; observation space is the set of possible responses to queries R . The dynamics of the system is simplified by the fact that the state transition function is trivial:

the underlying utility functions never change throughout the interaction process; the observation function is the response model which maintains a probability distribution of a particular response to a given query for a specific utility function; and, the reward function simply assigns costs to queries and expected utilities to final decisions.

Solving the preference elicitation POMDP is a difficult task. In realistic situations, the state space is continuous and multi-dimensional, so standard methods for solving finite-state POMDPs are no longer applicable. Boutilier [2002] presents a value-iteration based method that exploits the special structure inherent in the preference elicitation process to deal with parameterized belief states over the continuous state space; belief states are represented by truncated Gaussian or uniform mixture models. With standard gamble comparison queries that “slice” the density vertically (“Is utility of outcome x greater than p ?), updated distributions remain conjugate to the prior. The POMDP is solved by approximating the value function using asynchronous value iteration.

The preference POMDP can also be solved using policy-based methods. Braziunas and Boutilier [2004] describe an algorithm BBSLS that performs stochastic local search in the space of finite state policy controllers. In the case of continuous utility functions, it is possible to *sample* a number of states (utility functions) at each step, and calculate the observation and reward functions for the sampled states. The results for a very small preference elicitation problem from [Boutilier, 2002] provide the proof-of-concept verification of the policy-based approach. There is a lot of room for future research in this area as POMDP-based methods so far can only solve unrealistically small problems.

References

- Ali Abbas. Entropy methods for adaptive utility elicitation. *IEEE Transactions on Systems, Science and Cybernetics*, 34(2):169–178, 2004.
- G. Anandalingam and Chelsea C. White. A penalty function approach to alternative pairwise comparisons in ISMAUT. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(1):330–333, 1993.
- Fahiem Bacchus and Adam Grove. Graphical models for preference and utility. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence (UAI-95)*, pages 3–10, Montreal, 1995.
- Mokhtar S. Bazaraa and C. M. Shetty. *Nonlinear Programming: Theory and Algorithms*. Wiley, New York, 1979.
- Board on Manufacturing and Engineering Design. *Theoretical Foundations for Decision Making in Engineering Design*. National Academies Press, 2001.
- Craig Boutilier. A POMDP formulation of preference elicitation problems. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI-02)*, pages 239–246, Edmonton, 2002.
- Craig Boutilier. On the foundations of expected utility. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03)*, pages 285–290, Acapulco, 2003.
- Craig Boutilier, Fahiem Bacchus, and Ronen I. Brafman. UCP-Networks: A directed graphical representation of conditional utilities. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence (UAI-01)*, pages 56–64, Seattle, 2001.
- Craig Boutilier, Ronen Brafman, Chris Geib, and David Poole. A constraint-based approach to preference elicitation and decision making. In *AAAI Spring Symposium on Qualitative Decision Theory*, Stanford, 1997.
- Craig Boutilier, Ronen I. Brafman, Carmel Domshlak, Holger H. Hoos, and David Poole. Preference-based constrained optimization with CP-nets. *Computational Intelligence*, 20(2):137–157, 2004a.
- Craig Boutilier, Ronen I. Brafman, Holger H. Hoos, and David Poole. Reasoning with conditional ceteris paribus preference statements. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence (UAI-99)*, pages 71–80, Stockholm, 1999.
- Craig Boutilier, Rajarshi Das, Jeffrey O. Kephart, Gerald Tesauro, and William E. Walsh. Cooperative negotiation in autonomic systems using incremental utility elicitation. In *Proceedings of the Nineteenth Conference on Uncertainty in Artificial Intelligence (UAI-03)*, pages 89–97, Acapulco, 2003a.
- Craig Boutilier, Relu Patrascu, Pascal Poupart, and Dale Schuurmans. Constraint-based optimization with the minimax decision criterion. In *Ninth International Conference on Principles and Practice of Constraint Programming*, pages 168–182, Kinsale, Ireland, 2003b.
- Craig Boutilier, Relu Patrascu, Pascal Poupart, and Dale Schuurmans. Regret-based utility elicitation in constraint-based decision problems. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI-05)*, Edinburgh, 2005.
- Craig Boutilier, Tuomas Sandholm, and Rob Shields. Eliciting bid taker non-price preferences in (combinatorial) auctions. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI-04)*, pages 204–211, San Jose, CA, 2004b.

- Craig Boutilier, Richard S. Zemel, and Benjamin Marlin. Active collaborative filtering. In *Proceedings of the Nineteenth Conference on Uncertainty in Artificial Intelligence (UAI-03)*, pages 98–106, Acapulco, 2003c.
- Ronen Brafman and Carmel Domshlak. Introducing variable importance tradeoffs into CP-nets. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI-02)*, pages 69–76, Edmonton, 2002.
- Darius Braziunas and Craig Boutilier. Stochastic local search for pomdp controllers. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI-04)*, pages 690–696, San Jose, CA, 2004.
- Darius Braziunas and Craig Boutilier. Local utility elicitation in GAI models. In *Proceedings of the Twenty-first Conference on Uncertainty in Artificial Intelligence (UAI-05)*, pages 42–49, Edinburgh, 2005.
- Jack S. Breese, David Heckerman, and Carl Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence (UAI-98)*, pages 43–52, Madison, WI, 1998.
- David C. Brown and William P. Birmingham. Understanding the nature of design. *IEEE Expert*, 12(2): 14–16, 1997.
- Urszula Chajewska, Lise Getoor, Joseph Norman, and Yuval Shahar. Utility elicitation as a classification problem. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence (UAI-98)*, pages 79–88, Madison, WI, 1998.
- Urszula Chajewska and Daphne Koller. Utilities as random variables: Density estimation and structure discovery. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence (UAI-00)*, pages 63–71, Stanford, 2000.
- Urszula Chajewska, Daphne Koller, and Dirk Ormoneit. Learning an agent’s utility function by observing behavior. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML-01)*, pages 35–42, 2001.
- Urszula Chajewska, Daphne Koller, and Ronald Parr. Making rational decisions using adaptive utility elicitation. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-00)*, pages 363–369, Austin, TX, 2000.
- William W. Cohen, Robert E. Shapire, and Yoram Singer. Learning to order things. *Journal of Artificial Intelligence Research*, 1999.
- Richard M. Cyert and Morris H. de Groot. Adaptive utility. In Maurice Allais and Ole Hagen, editors, *Expected Utility Hypothesis and the Allais Paradox*, pages 223–241. D. Reidel, Dordrecht, 1979.
- Joseph G. D’Ambrosio and William P. Birmingham. Preference-directed design. *Journal for Artificial Intelligence in Engineering Design, Analysis and Manufacturing*, 9:219–230, 1995.
- Morris H. de Groot. Decision making with an uncertain utility function. In Bernt P. Stigum and Fred Wenstøp, editors, *Foundations of Utility and Risk Theory with Applications*, pages 371–384. D. Reidel, Dordrecht, 1983.
- Rina Dechter. *Constraint Processing*. Morgan Kaufmann, San Francisco, 2003.

- Jon Doyle, Yoav Shoham, and Michael P. Wellman. A logic of relative desire. In Z. W. Ras and M. Zemanekova, editors, *Methodologies for Intelligent Systems 6*, pages 16–31. Springer-Verlag, Berlin, 1991.
- Jon Doyle and Richmond H. Thomason. Background to qualitative decision theory. *AI Magazine*, 20(2): 55–68, 1999.
- Jon Doyle and Michael P. Wellman. Representing preferences as *ceteris paribus* comparatives. In *Proceedings of the AAAI Spring Symposium on Decision-Theoretic Planning*, pages 69–75, Stanford, 1994.
- James S. Dyer. Interactive goal programming. *Management Science*, 19:62–70, 1972.
- Peter H. Farquhar. Utility assessment methods. *Management Science*, 30(11):1283–1300, 1984.
- Peter C. Fishburn. *Decision and Value Theory*. Wiley, New York, 1964.
- Peter C. Fishburn. Interdependence and additivity in multivariate, unidimensional expected utility theory. *International Economic Review*, 8:335–342, 1967a.
- Peter C. Fishburn. Methods of estimating additive utilities. *Management Science*, 13:435–453, 1967b.
- Peter C. Fishburn. *Utility Theory for Decision Making*. Wiley, New York, 1970.
- Peter C. Fishburn. Multiattribute utilities in expected utility theory. In David E. Bell, Ralph L. Keeney, and Howard Raiffa, editors, *Conflicting Objective in Decisions*, pages 172–196. Wiley, 1977.
- Peter C. Fishburn. *The Foundations of Expected Utility*. D. Reidel, Dordrecht, 1982.
- Simon French. *Decision Theory*. Halsted Press, New York, 1986.
- Simon French and David Ríos Insua. *Statistical decision theory*. Oxford University Press, New York, 2000.
- Soumyadip Ghosh and Jayant Kalagnanam. Polyhedral sampling for multiattribute preference elicitation. In *Proceedings of the Fifth ACM Conference on Electronic Commerce*, pages 256–257, San Diego, CA, 2003.
- Ken Goldberg, Theresa Roeder, Dhruv Gupta, and Chris Perkins. Eigentaste: A constant time collaborative filtering algorithm. *Information Retrieval*, 4(1):133–151, 2001.
- Christophe Gonzales and Patrice Perny. GAI networks for utility elicitation. In *Proceedings of the Ninth International Conference on Principles of Knowledge Representation and Reasoning (KR2004)*, pages 224–234, Whistler, BC, 2004.
- Paul E. Green and Vithala R. Rao. Conjoint measurement for quantifying judgmental data. *Journal of Marketing Research*, 8(3):355–363, August 1971.
- Paul E. Green and V. Srinivasan. Conjoint analysis in consumer research: Issues and outlook. *Journal of Consumer Research*, 5(2):103–123, September 1978.
- Paul E. Green and V. Srinivasan. Conjoint analysis in marketing: New developments with implications for research and practice. *Journal of Marketing*, 54(4):3–19, October 1990.
- Vu Ha and Peter Haddawy. Toward case-based preference elicitation: Similarity measures on preference structures. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence (UAI-98)*, pages 193–201, Madison, WI, 1998.

- John C. Harsanyi. Games with incomplete information played by Bayesian players. part I: The basic model. *Management Science*, 14:159–182, 1967.
- John C. Harsanyi. Games with incomplete information played by Bayesian players. part II: Bayesian equilibrium points. *Management Science*, 14:320–334, 1968.
- Gordon B. Hazen. Partial information, dominance, and potential optimality in multiattribute utility theory. *Operations Research*, 34(2):296–310, 1986.
- Hillary A. Holloway and Chelsea C. White, III. Question selection for multiattribute decision-aiding. *European Journal of Operational Research*, 148:525–543, 2003.
- Ronald A. Howard and James E. Matheson, editors. *Readings on the Principles and Applications of Decision Analysis*. Strategic Decision Group, Menlo Park, CA, 1984.
- Vijay S. Iyengar, Jon Lee, and Murray Campbell. Q-Eval: Evaluating multiple attribute items using queries. In *Proceedings of the Third ACM Conference on Electronic Commerce*, pages 144–153, Tampa, FL, 2001.
- Ralph L. Keeney and Howard Raiffa. *Decisions with Multiple Objectives: Preferences and Value Trade-offs*. Wiley, New York, 1976.
- Craig W. Kirkwood and Rakesh K. Sarin. Ranking with partial information: A method and an application. *Operations Research*, 33(1):38–48, 1985.
- Joseph A. Konstan, Bradley N. Miller, David Maltz, Jonathan L. Herlocker, Lee R. Gordon, and John Riedl. GroupLens: Applying collaborative filtering to usenet news. *Communications of the ACM*, 40(3):77–87, 1997.
- Risto Lahdelma, Joonas Hokkanen, and Pekka Salminen. SMAA - stochastic multiobjective acceptability analysis. *European Journal of Operational Research*, 106:137–143, 1998.
- Andreu Mas-Colell, Micheal D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
- Michael McGeachie and Jon Doyle. Utility functions for ceteris paribus preferences. *Computational Intelligence*, 20(2), 2004.
- Andrew Ng and Stuart Russell. Algorithms for inverse reinforcement learning. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML-00)*, pages 663–670, Stanford, CA, 2000.
- Relu Patrascu, Craig Boutilier, Rajarshi Das, Jeffrey O. Kephart, Gerald Tesauro, and William E. Walsh. New approaches to optimization and utility elicitation in autonomic computing. In *Proceedings of the Twentieth National Conference on Artificial Intelligence*, 2005.
- Steve Prestwich, Francesca Rossi, Brent Venable, and Toby Walsh. Constraint-based preferential optimization. In *Proceedings of the Twentieth National Conference on Artificial Intelligence*, Pittsburgh, 2005.
- Pearl Pu, Boi Faltings, and Marc Torrens. User-involved preference elicitation. In *IJCAI-03 Workshop on Configuration*, Acapulco, 2003.

- Thomas L. Saaty. A scaling method for priorities in hierarchical structures. *Journal of Mathematical Psychology*, 15:234–281, 1977.
- Thomas L. Saaty. *The analytic hierarchy process: planning, setting priorities, resource allocation*. McGraw-Hill, 1980.
- Ahti Salo and Raimo P. Hämäläinen. Preference programming through approximate ratio comparisons. *European Journal of Operational Research*, 82:458–475, 1995.
- Ahti Salo and Raimo P. Hämäläinen. Preference ratios in multiattribute evaluation (PRIME)–elicitation and decision procedures under incomplete information. *IEEE Trans. on Systems, Man and Cybernetics*, 31(6):533–545, 2001.
- Ahti Salo and Raimo P. Hämäläinen. Preference programming. Manuscript available at <http://www.sal.hut.fi/Publications/pdf-files/msal03b.pdf>, 2004.
- Ahti A. Salo and Raimo P. Hämäläinen. On the measurement of preferences in the Analytic Hierarchy Process. *Journal of Multi-Criteria Decision Analysis*, 6:309–319, 1997.
- Rakesh K. Sarin. Screening of multiattribute alternatives. *Omega*, 5(4):481–489, 1977.
- Leonard J. Savage. The theory of statistical decision. *Journal of the American Statistical Association*, 46(253):55–67, 1951.
- Leonard J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.
- Edward A. Sykes and Chelsea C. White, III. Multiobjective intelligent computer-aided design. *IEEE Transactions on Systems, Man and Cybernetics*, 21(6):1498–1511, 1991.
- Olivier Toubia, John R. Hauser, and Duncan I. Simester. Polyhedral methods for adaptive choice-based conjoint analysis. *Journal of Marketing Research*, 41(1):116–131, 2004.
- Olivier Toubia, Duncan I. Simester, John R. Hauser, and Ely Dahan. Fast polyhedral adaptive conjoint estimation. *Marketing Science*, 22(3):273–303, 2003.
- John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, 2nd edition, 1947.
- A. Wald. *Statistical Decision Functions*. Wiley, New York, 1950.
- Tianhan Wang and Craig Boutilier. Incremental utility elicitation with the minimax regret decision criterion. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03)*, pages 309–316, Acapulco, 2003.
- Martin Weber. Decision making with incomplete information. *European Journal of Operational Research*, 28:44–57, 1987.
- Chelsea C. White, Shigeru Dozono, and William T Scherer. An interactive procedure for aiding multiattribute alternative selection. *Omega*, 11(2):212–214, 1983.
- Chelsea C. White, III, Andrew P. Sage, and Shigeru Dozono. A model of multiattribute decisionmaking and trade-off weight determination under uncertainty. *IEEE Transactions on Systems, Man and Cybernetics*, 14(2):223–229, 1984.
- Bob Wielinga and Guus Schreiber. Configuration-design problem solving. *IEEE Expert*, 12(2):49–56, 1997.