Reprinted from *Proceedings of the First International Conference on Artificial Intelligence Planning Systems (AIPS-92)*, June, 1992, pp. 236–242.

Modular Utility Representation for Decision-Theoretic Planning

Michael P. Wellman
USAF Wright Laboratory
WL/AAA-1
Wright-Patterson AFB, OH 45433
(513) 255-5800
wellman@wl.wpafb.af.mil

Jon Doyle
MIT Laboratory for Computer Science
545 Technology Square
Cambridge, MA 02139
(617) 253-3512
doyle@zermatt.lcs.mit.edu

Abstract

Specification of objectives constitutes a central issue in knowledge representation for planning. Decision-theoretic approaches require that representations of objectives possess a firm semantics in terms of utility functions, yet provide the flexible compositionality needed for practical preference modeling for planning systems. Modularity, or separability in specification, is the key representational feature enabling this flexibility. In the context of utility specification, modularity corresponds exactly to well-known independence concepts from multiattribute utility theory, and leads directly to approaches for composing separate preference specifications. Ultimately, we seek to use this utility-theoretic account to justify and improve existing mechanisms for specification of preference information, and to develop new representations exhibiting tractable specification and flexible composition of preference criteria.

1 REPRESENTING UTILITY FOR PLANNING

As generally conceived, the AI planning task aims to use beliefs about the world and predicted effects of available actions to synthesize a course of action furthering some objectives. *Decision-theoretic* planning, which measures beliefs in terms of Bayesian probability and objectives in terms of expected utility, challenges the architect of planning systems to design representation constructs that can be interpreted faithfully in terms of probabilities and utilities, that can be scaled to facilitate expression of general knowledge about broad domains, and that support computationally tractable inference about plans and partial plans.

Multiple objectives, partially achievable objectives, and uncertainty about the effects of actions all pose difficult problems for traditional goal-based planning systems. In themselves, goal conditions provide no means for resolving tradeoffs among competing objectives or for expressing varying degrees of partial satisfaction. Since these difficulties arise in most (if not all) realistic planning problems, builders of practical planning systems commonly augment goal-based representations with heuristic measures of goal importance, partial goal achievement, or other priority relationships. These augmentations might assign, for example, numeric achievement values to individual goals, costs to individual actions, and penalties proportional to the measured distance from a goal. The planner then combines them in some straightforward way (e.g., achievement values minus costs minus penalties) to evaluate an overall plan. Although such *ad hoc* mechanisms might provide reasonable performance in particular planning systems, they typically lack any precise meaning, and so provide neither a basis for evaluating their coherence and appropriateness for other problem domains, nor a justification for the inference operations and choices executed by the underlying planning architecture.

In contrast, interpreting specifications of objectives in terms of decision-theoretic preferences permits the designers of planning systems to judge both the coherence of the objectives and the effectiveness of the planning system in furthering these objectives. But decision theory per se does not address the problem of designing convenient representations for preference, or its corresponding measure, utility. Applying the concepts of decision theory directly (as in *decision analysis*) requires specifying a utility function over the entire domain, ranking plan results by their desirability in any conceivable planning situation. This places unwarranted burdens on the modeler, since different features of the situation are relevant with respect to decisions made at different levels of abstraction or at intermediate stages of plan synthesis. To make utility specifications more convenient, we seek *modular* representations that separately specify preference information concerning particular factors, so that we can dynamically combine those factors deemed relevant to a particular problem and level of abstraction.

The following builds a framework for modular specification of utilities on firm decision-theoretic foundations. We begin by presenting our view of modularity in knowledge representation as specification of flexibly composable model elements. Next, we present some background material on multiattribute utility theory prerequisite to our account of modular utility specification. We then demonstrate the correspondence between separability in specification and well-established independence concepts from utility theory, and exhibit the consequences of the theory for composition operations on utility representations. We conclude with a summary discussion of related issues and work.

2 MODULARITY AND MODEL COMPOSITION

AI planning distinguishes itself from other approaches to automated decision making by emphasizing compositional synthesis of a course of action from primitive action elements (i.e., operators) together with specifications of the effects of each of these primitive elements in isolation. To synthesize a composite plan, a planner must determine the overall effects of the composite plan as a modular combination of the effects of its constituent actions. Modular specification of effects is essential to giving planners the freedom to compose primitives as necessary. But modular specification of planning *objectives* is equally important. Access to preferences regarding specific outcome features (without specifying them over complete outcomes) is essential when it is impossible or infeasible to characterize the entire outcome space in advance, when different features are relevant for different decision problems, and when preferences for particular aspects of the outcome depend on background context.

For example, consider the problem of planning large-scale military transportation operations. At a high level, we might consider monetary costs and whether the specified movement requirements are met, whereas a more detailed analysis would consider the timeliness of cargo movements, the amount and type of cargo moved, stress on transportation resources, and safety. When making isolated decisions about parts of the operation (the usual case), it often proves advantageous to treat resource reservations that impact the rest of the plan as part of the outcome, and to summarize the value of those resources as opportunity costs.

Assessing a global utility function covering all of these outcome features and their subconcepts seems impractical. We believe it more reasonable to specify utility functions over individual features or small groups of features, combining these as needed for making tradeoffs in decision problems involving sets of features. For example, we might have a measure of the value of moving various types of cargo and a relation describing the tradeoff between monetary cost and tardiness for particular movement classes. When faced with a particular decision problem, the planner assembles the relevant outcome features and corresponding utility specifications, then constructs a comprehensive utility model by composing the individual utilities.[1]

To realize this approach, we must develop interpretations for isolated preference specifications, and methods for defining composition operators and composing selected partial specifications. Fortunately, utility theorists have developed a rich framework for analyzing the composition of utility functions over multiple attributes, motivated by the need to simplify assessment even when attributes are fixed in advance.[2] We can exploit this theory for more flexible utility representation as well, both to make sense of modular specifications and to determine the appropriate form of composition operators. In the remainder of this paper, we present the relevant utility theory, and demonstrate its application to the problem of flexible composition of modular utility specifications.

3 PREFERENCES AND UTILITY FUNCTIONS

Utility theory starts with the notion of *preferences* over *outcomes* (Keeney and Raiffa, 1976). Outcomes represent the possible consequences of the agent's decisions. In the planning context, an outcome might be taken to be the state resulting from execution of a plan, or perhaps an entire history of instantaneous states over the lifetime of the agent. To provide an adequate basis for decision, the set Ω of (mutually incompatible) possible outcomes must distinguish all consequences that the agent cares about and are possibly affected by its actions. We define the agent's preferences by a total preorder (a complete, reflexive, and transitive relation), $\succsim$, over possible outcomes, called the *preference order*.

Given a few topological restrictions on $\succsim$, the preference order can be captured by an order-preserving, real-valued *utility function*, u. The function u *represents* $\succsim$ in the sense that outcomes can be ranked by comparing the numeric values of the utility function

[1]The task of dynamically constructing decision models for particular problem instances from general domain knowledge is gaining increasing attention from researchers (Wellman *et al.*, 1992). Most of the work to date concerns probabilistic modeling, but some addresses the equally important problem of generating utility models (Haddawy and Hanks, 1990; Loui, 1990).

[2]Researchers in multiattribute utility theory tend to refer to *de*composition rather than composition because they look to assess a fixed outcome space in a top-down manner, rather than the bottom-up assembly of primitive utility specifications.

applied to those outcomes, with $\omega \succsim \omega'$ iff $u(\omega) \geq u(\omega')$ for $\omega, \omega' \in \Omega$. However, the utility-function representation is not unique. If $u(\omega) \geq u(\omega')$, then it must also be the case that $\varphi(u(\omega)) \geq \varphi(u(\omega'))$, for any monotonically increasing function φ. Since they represent the same preference order—and thus would sanction identical decisions (under certainty)—we say that u and $\varphi \circ u$ are *strategically equivalent*.

When there is *uncertainty*, plans influence outcomes only probabilistically, and we must represent the anticipated result of a plan by a probability distribution over Ω, or *prospect*, and extend $\succsim$ to order prospects. The central result of utility theory is a representation theorem that establishes (given some restrictions on $\succsim$) the existence of a utility function $u : \Omega \to \mathbf{R}$ such that preference over prospects is represented by the *expectation* of u over those prospects. The key point here is that u is defined over outcomes alone; the extension to prospects via expectation is a consequence of the axioms of probability and utility (Savage, 1972).

As in the deterministic case, the utility-function representation for a preference order over uncertain prospects is not unique. However, monotone transformations do not generally preserve expectation orderings, and hence the class of strategically equivalent functions is more limited in the uncertain case. Specifically, expected utility functions[3] are unique up to a positive linear transformation. That is, for positive linear function ψ (i.e., $\psi(x) = ax + b$, $a > 0$), the utility functions u and $\psi \circ u$ are strategically equivalent.

In principle, we could avoid utility functions altogether and perform decision-theoretic reasoning directly in terms of preference orders, but numeric representations offer distinct advantages in compactness and analytic manipulability. Multiattribute utility theory exploits these advantages as far as possible, by decomposing complex outcome spaces into modular structures and specifying complex utility functions in terms of combinations of lower-dimension functions.

4 MULTIATTRIBUTE OUTCOMES

In realistic decision situations, there are commonly many objectives, and hence an outcome would represent a highly complex set of features describing the plan's result. In utility theory, preference-relevant features of an outcome are called *attributes*, and a substantial body of work has been devoted to relating preference on individual attributes to preference for complex outcomes. The first requirement for separating overall preference into that for individual attributes is some structure on the outcome space. This structure is provided by a *framing*, which defines a multiattribute representation of the given outcome space (Wellman and Doyle, 1991). A framing represents each outcome ω as a vector $\langle \omega_1, \ldots, \omega_n \rangle$ of attributes, where each outcome attribute ω_i is drawn from an *attribute space* A_i. The framing thus views the set of outcomes Ω as a subset of the multiattribute space $A = \prod_{i=1}^{n} A_i$. In the following, we assume that every attribute vector corresponds to some outcome, that is, that the attributes characterize the outcomes exactly.[4]

Given a particular framing, the utility function over Ω can be expressed as a multi-dimensional function over A. The modularity advantages of the multiattribute representation accrue if this multi-dimensional function can in turn be decomposed into some regular combination of lower dimension *subutility* functions, each representing preferences over one or more attributes.

5 SEPARABILITY AND UTILITY INDEPENDENCE

Indirect specification of a multi-dimensional function as a combination of functions of lower dimension depends on separability of the various dimensions. In the context of utility theory, we are further concerned that the lower-dimension functions themselves have some meaningful interpretation in terms of preferences, ideally that they be considered subutility functions in some sense. What, then, does it mean to say that u_i is a subutility function for attribute i? We can answer this question in terms of invariance of decisions, or strategic equivalence.

If we wish to interpret $u_i : A_i \to \mathbf{R}$ as a utility function, then it must correspond to some preference order, $\succsim_i$, over prospects involving A_i. To talk sensibly about preferences over A_i without referring to the remainder of the outcome vector, preferences for attribute i must be invariant in some sense with respect to the other attributes. If we indeed specify an entire utility function for that attribute, we are uniquely determining $\succsim_i$, and in effect determining all decisions involving that attribute, assuming that all others are fixed. In other words, u_i determines the optimal decision (or decisions, in case of ties) for all choices involving prospects where outcomes differ only in attribute i. Moreover, this decision does not depend on what fixed values the other attributes take.

This invariance property is a fundamental concept in multiattribute utility theory, called *utility independence* (UI) (Keeney and Raiffa, 1976).

[3]Henceforth, we assume the more general, uncertain, case, and refer exclusively to utility functions that exhibit the expectation property.

[4]For our purposes here (although not for other purposes (Wellman and Doyle, 1991)), we can safely satisfy this assumption by padding the outcome space and extending the preference order in a manner consistent with other given constraints.

Definition 1 (UI) *One attribute is* utility independent *(UI)*[5] *of the remaining attributes if preferences for prospects over this attribute, holding other attribute values fixed, do not depend on the fixed values of those attributes.*

We can generalize this to UI between two sets of attributes by considering prospects where attributes in the first set vary and those in the second set are fixed. Note that UI is not symmetric—for instance, when $n = 3$, it is quite possible that $\{A_1\}$ be UI of $\{A_2, A_3\}$, but $\{A_2, A_3\}$ not be UI of $\{A_1\}$.

Utility independence appears to be a minimal requirement for modular specification of preferences. The reason is that without implicitly invoking UI relationships, it is generally not coherent to refer to preferences on individual outcome features via subutility functions. If we accept the view that specifying properties of subutility functions on individual features or small groups of features is the essence of modularity, we find that extensive application of utility independence is pragmatically unavoidable.

Fortunately, the UI condition justifies some strong separability results, leading to well-structured utility representations. The separable form of a utility function over UI attributes follows directly from the invariance condition. Suppose $n = 2$, and A_1 is UI of A_2. The overall utility is a function of both attributes, $u : A_1 \times A_2 \to \mathbf{R}$. We know from the UI condition that the conditional utility function $u(\cdot, \omega_2')$ where the second attribute is fixed at the constant value ω_2' must be strategically equivalent (with respect to the first attribute) to the conditional utility function corresponding to any other value. Because utility is unique up to a positive linear transformation, this implies, for utility conditioned on ω_2'', that

$$u(\omega_1, \omega_2'') = au(\omega_1, \omega_2') + b$$

for some constants $a > 0$ and b. Indeed, such a relationship must hold for any value $\omega_2 \in A_2$, although the a and b parameters may depend on ω_2. This observation yields a general form for the overall utility function,

$$u(\omega_1, \omega_2) = g(\omega_2)u(\omega_1, \omega_2') + h(\omega_2).$$

Defining the subutility function $u_1(\omega_1)$ as $u(\omega_1, \omega_2')$, for the particular constant ω_2', we have

$$u(\omega_1, \omega_2) = g(\omega_2)u_1(\omega_1) + h(\omega_2). \quad (1)$$

In fact, the functions g and h can also be expressed in terms of conditional utility functions for attribute 2, with the first attribute fixed at particular values. (See Keeney and Raiffa (1976, Chapter 5) for the details of this decomposition.)

[5]We use the same abbreviation for both noun and adjective forms of the concept.

6 MULTIATTRIBUTE UTILITY FORMS

The development above establishes the separability of an attribute subset from its complement in the framing, when UI holds. In framings with many attributes, we would expect to have partial preference information in the form of subutility functions (hence implicit UI conditions) corresponding to a variety of attribute subsets and individual attributes. When certain patterns of UI hold over the entire multiattribute space, general forms for the overall utility function follow. There are two elementary forms of multiattribute decomposition, in both of which the overall utility function can be expressed as a modular combination of single-attribute subutility functions.[6]

In the *multilinear* decomposition, the n-dimensional utility function is separable into $n-1$ subutility functions for individual attributes, with perhaps one other (which we take to be the first, without loss of generality) not expressible as a subutility,

$$u(\omega_1, \ldots, \omega_n) = f(\omega_1, u_2(\omega_2), \ldots, u_n(\omega_n)).$$

The form is called multilinear because the function f is linear in each argument save the first, holding the remaining arguments fixed. It is a valid decomposition as long as each attribute (except possibly the first) is UI of the rest, that is, has preferences validly expressed as a subutility function. The disadvantage of the multilinear form is that it requires that we specify $O(2^n)$ parameters (called *scaling constants*) in addition to the single-attribute functions.

The *multiplicative* form corresponds to the sum or product of the subutility functions, each weighted by a scaling constant. Since this form requires only $O(n)$ parameters, it is far easier to specify (i.e., more modular) than the multilinear form. The price paid for this simplicity is that each subset of the attributes must be UI of its complement. Of course, we could not expect to have explicit UI assertions or subutilities corresponding to all subsets; specifying these would defeat the purpose of modularity anyway. Fortunately, the utility independence of some sets of attributes often entails the UI of related sets. The theory of UI relations provides a basis for deriving the most modular form corresponding to a given set of fundamental independence relations.

The basic mechanism for deriving new separability conditions from a set of UI relations is based on a result originally due to Gorman (1968). Suppose we have two attribute sets, each UI of its complement,

[6]The decompositions actually require a slightly weaker condition than UI, called *generalized utility independence* (GUI) (Keeney and Raiffa, 1976) or *autonomy* (von Stengel, 1988). In the following discussion, we continue to appeal to UI, since the preferential interpretation of subutility functions implicitly invokes the stronger concept.

with a nonempty intersection, Y. We can write the first set as $X \cup Y$ and the second as $Y \cup Z$, with X, Y, and Z disjoint attribute sets. Then it follows that all combinations of these sets—X, Y, Z, $X \cup Z$, and $X \cup Y \cup Z$—are also UI[7] of their respective complements. From this fact, we see that a small number of UI conditions for overlapping attribute sets can implicitly entail a large number of independence relations.

7 MODULAR COMPOSITION OF UTILITY FUNCTIONS

The general utility-model composition problem is as follows. Suppose we are given a collection of subutility functions and other preference information involving attributes from a given framing.[8] Taking the existence of subutility functions to implicitly assert a corresponding UI condition, our task is to find a modular composition of these subutilities into an overall utility function. In doing so, we exploit the UI conditions entailed by those implicit in the given subutility functions according to the rule described above.

First, note that determining the form of the overall utility function is not simply a matter of verifying whether the given UI conditions collectively justify an n-attribute multiplicative or multilinear decomposition. There is actually a structural continuum between these forms, defined by the space of hierarchical decompositions in which each node is a multiplicative or multilinear function of some partition of the attribute set. For example, with $n = 5$, the top-level decomposition might be a multilinear combination of the form

$$u(\omega_1, \ldots, \omega_5) = f(\omega_1, u_2(\omega_2), u_{3,4,5}(\omega_3, \omega_4, \omega_5)),$$

and $u_{3,4,5}$ might be recursively decomposable as a multiplicative combination of its three attributes.

In fact, there exists a unique decomposition hierarchy, or *utility tree*, corresponding to any set of UI conditions (Gorman, 1968; von Stengel, 1988). Moreover, we can derive this tree from a given set of UI premises without enumerating all of the UI relationships that follow from these premises. We have developed an algorithm (to be reported in detail elsewhere) that computes the decomposition hierarchy corresponding to an arbitrary collection of UI assertions in an incremental manner, permitting a "structural sensitivity analysis" of the implications of additional UI axioms or subutility functions. The basic ideas in the algorithm follow from the constructive demonstration of the decomposition theorems by Gorman (1968), Keeney and Raiffa (1976), and von Stengel (1988).

[7]Technically, only in the generalized sense noted above.

[8]A more flexible approach would be to construct the framing dynamically based on attributes for which we have preferential information. In this discussion, we abstract from the problem of defining the framing, as well as that of obtaining the subutility functions, although both are critical steps in utility model composition.

8 AN EXAMPLE

To illustrate the hierarchical structure of a multiattribute utility function, we adduce an example from the transportation domain. The attributes and independence relations we employ are selected purely for expository purposes and are not intended to represent real salience or independence in this domain.

Suppose we are considering alternate modes of transportation for a particular cargo movement, say M. To evaluate the results of this movement, we might consider the following outcome attributes: (1) amount of bulk cargo transported in M, (2) monetary expenses associated with M, (3) tardiness of M with respect to some target arrival time, (4) opportunity costs associated with vehicles employed in M, (5) opportunity costs of facilities (e.g., warehouses, loading equipment) employed in M, (6) human resources (e.g., vehicle crews) used in M, and (7) safety. In a more concrete instance, these attributes would be directly associated with particular resource and cargo types.

Suppose further that we have preference information about these attributes in the form of subutility functions. Let $u_{i,j,\ldots}$ denote the subutility function corresponding to attributes $i, j, \ldots$ in the numbering above. In our example, suppose that we have (given or perhaps derived from some more fundamental information) the subutility functions:

$$u_1,\ u_{2,3},\ u_{2,4,5,6},\ u_{4,5},\ \text{and}\ u_7.$$

That is, we have subutility functions for the individual attributes bulk cargo movement and safety, and joint subutility functions describing the tradeoff between monetary expenses and tardiness, as well as among monetary expenses and the various resource costs. We also have a specification of the particular tradeoff between vehicle and facility resource usage.

We interpret the existence of these subutility functions as implicitly asserting UI between the domain of each function and the rest of the attributes. These UI conditions, in turn, lead to a unique utility tree describing the modular composition of these subutilities into an overall utility function. Figure 1 depicts the tree corresponding to the subutility functions listed above.

At the top level, the utility function is a multilinear combination of the given subutilities u_1 and u_7, along with a joint subutility function for the remaining attributes. This in turn is composed of subutilities u_2, u_3, and $u_{4,5,6}$, none of which are among those originally specified. However, all of the leaf subutilities are derivable from the originals given the implied UI conditions.[9] For example, u_2 and u_3 are condi-

[9]As are many of the necessary scaling constants.

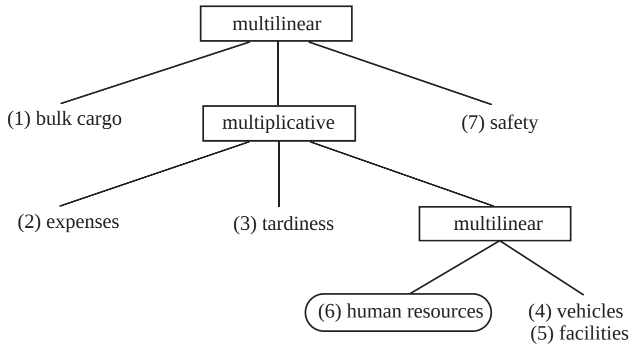

Figure 1: The utility tree for the transportation example. Except for *human resources*, each of the leaves corresponds to a separable subutility function.

tioned versions of $u_{2,3}$, obtained by fixing attributes 3 and 2, respectively, at arbitrary levels. These are indeed subutilities—hence the freedom in choosing conditioning values—by virtue of the UI relations implicit in the overlapping UI index sets $\{2, 3\}$ and $\{2, 4, 5, 6\}$. If in fact $u_{2,3}$ is not separable, then this subutility function is incompatible with the existence of subutility $u_{2,4,5,6}$.[10] Finally, the subutility $u_{4,5,6}$ is separable via a two-attribute multilinear form. This composition corresponds to the UI form (Equation (1)) separating $\{4, 5\}$ from $\{6\}$. Note that $\{6\}$ is not UI of $\{4, 5\}$, and therefore preference for attribute 6 is not expressible as a subutility. And note also that attributes 4 and 5 are not separable; the utility tree employs the joint subutility $u_{4,5}$.

9 CONCLUSIONS

In summary, we have argued that modularity is an essential feature of utility representation for decision-theoretic planning, and that separate specification of utility for isolated outcome features is tantamount to an assertion of utility independence. The independence relations implicit in a collection of modular utility components dictate the form in which they should be composed to define an overall utility function. Customized utility models can be constructed dynamically to reflect the relevant factors in a particular decision situation via bottom-up composition according to rules of multiattribute utility theory.

When the modular preference specifications are partial (i.e., not complete subutility functions), the corresponding invariance property implicit in the separation is generally weaker than utility independence. For example, specifying only ordinal subutilities (e.g., monotonicity conditions) is tantamount to *preferential independence*, another well-known concept of utility theory. In principle, we could extend this analysis to cover a broad spectrum of independence concepts with a variety of implications for utility composition.

Our investigation is in the spirit of analogous analyses relating modularity and probabilistic independence (Heckerman, 1990; Heckerman and Horvitz, 1988; Wellman, 1990). In both cases, we must invoke independence to justify scalable representation schemes, and we may exploit the independence relations to define valid composition rules. Utility trees can be viewed as an analog of probabilistic dependence graphs, the underlying basis for the most prevalent modeling scheme for probabilistic reasoning (Charniak, 1991; Pearl, 1988).

The notion of incremental utility specification and combination has also been advocated by Loui (1989; 1990). Our analysis serves to determine when such combination is sanctioned by utility theory, and to constrain the form that it might take. It may also be reasonable to heuristically apply modular combination rules when they are not theoretically justified, as the computational benefits may outweigh the cost of potential errors.

We are currently attempting to incorporate these ideas in the design of a scheme for utility representation to be used as part of common Knowledge Representation Specification Language for the DARPA/Rome Laboratory initiative on Transportation Planning. The experience gained in using these techniques to build a substantial KB for use by several research groups will be invaluable in developing more refined representations for decision-theoretic planning.

[10] The problem of dealing with overconstrained or otherwise inconsistent utility specifications is an important issue in utility composition but beyond the scope of this paper.

Acknowledgments

Our investigation of this topic arose from an effort to design a general utility representation for transportation planning. We have benefited from collaborations with the other members of this project: Jack Breese, Tom Dean, and Max Henrion. We are also grateful to Ron Loui and Bernhard von Stengel for past discussions about central aspects of this work. Jon Doyle is supported by the USAF Rome Laboratory and DARPA under contract F30602-91-C-0018, and by National Institutes of Health Grant No. R01 LM04493 from the National Library of Medicine.

References

(Charniak, 1991) Eugene Charniak. Bayesian networks without tears. *AI Magazine*, pages 50–63, Winter 1991.

(Gorman, 1968) W. M. Gorman. The structure of utility functions. *Review of Economic Studies*, 35:367–390, 1968.

(Haddawy and Hanks, 1990) Peter Haddawy and Steve Hanks. Issues in decision-theoretic planning:

Symbolic goals and numeric utilities. In *Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control*, pages 48–58, 1990.

(Heckerman and Horvitz, 1988) David E. Heckerman and Eric J. Horvitz. The myth of modularity in rule-based systems for reasoning with uncertainty. In John F. Lemmer and Laveen N. Kanal, editors, *Uncertainty in Artificial Intelligence 2*, pages 23–34. North-Holland, 1988.

(Heckerman, 1990) David E. Heckerman. *Probabilistic Similarity Networks*. PhD thesis, Stanford University, Stanford, CA, 1990.

(Keeney and Raiffa, 1976) Ralph L. Keeney and Howard Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley and Sons, New York, 1976.

(Loui, 1989) Ronald P. Loui. Defeasible decisions: What the proposal is and isn't. In *Proceedings of the Workshop on Uncertainty in Artificial Intelligence*, pages 245–252, Windsor, Ontario, 1989.

(Loui, 1990) Ronald Loui. Defeasible specification of utilities. In Henry E. Kyburg, Jr., Ronald P. Loui, and Greg N. Carlson, editors, *Knowledge Representation and Defeasible Reasoning*, pages 345–359. Kluwer Academic Publishers, 1990.

(Pearl, 1988) Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA, 1988.

(Savage, 1972) Leonard J. Savage. *The Foundations of Statistics*. Dover Publications, New York, second edition, 1972.

(von Stengel, 1988) Bernhard von Stengel. Decomposition of multiattribute expected-utility functions. *Annals of Operations Research*, 1988.

(Wellman and Doyle, 1991) Michael P. Wellman and Jon Doyle. Preferential semantics for goals. In *Proceedings of the National Conference on Artificial Intelligence*, pages 698–703, 1991.

(Wellman *et al.*, 1992) Michael P. Wellman, John S. Breese, and Robert P. Goldman. From knowledge bases to decision models. *The Knowledge Engineerting Review*, 7(1):35–53, 1992.

(Wellman, 1990) Michael P. Wellman. The STRIPS assumption for planning under uncertainty. In *Proceedings of the National Conference on Artificial Intelligence*, Boston, MA, 1990. American Association for Artificial Intelligence.