AI Magazine Volume 10 Number 1 (1989) (© AAAI)

Logical and Decision-Theoretic Methods for Planning under Uncertainty

Curtis P. Langlotz and Edward H. Shortliffe

Decision theory and nonmonotonic logics are formalisms that can be employed to represent and solve problems of planning under uncertainty. We analyze the usefulness of these two approaches by establishing a simple correspondence between the two formalisms. The analysis indicates that planning using nonmonotonic logic comprises two decision-theoretic concepts: probabilities (degrees of belief in planning hypotheses) and utilities (degrees of preference for planning outcomes). We present and discuss examples of the following lessons from this decision-theoretic view of nonmonotonic reasoning: (1) decision theory and nonmonotonic logics are intended to solve different components of the planning problem; (2) when considered in the context of planning under uncertainty, nonmonotonic logics do not retain the domain-independent characteristics of classical (monotonic) logic; and (3) because certain nonmonotonic programming paradigms (for example, frame-based inheritance, nonmonotonic logics) are inherently problem specific, they might be inappropriate for use in solving certain types of planning problems. We discuss how these conclusions affect several current AI research issues.

Planning consists of devising a course of action that conforms as well as possible to a set of goals. A planner attempts to determine the optimal action in a particular problem-solving situation. We are interested in automating decision support for a particular set of planning problems distinguished by the following characteristics: (1) the current situation is not known with certainty; (2) the consequences of action are not known with certainty; and (3) the goals of the planning process are conflicting and, therefore, are not completely satisfiable. We refer to problems of this type as planning under uncertainty. Because these planning tasks entail uncertainty and tradeoffs, a purely deductive process (such as state space search [Fikes and Nilsson 1971] or skeletal plan refinement [Friedland and Iwasaki 1985]) is difficult to employ to select the optimal plan (Langlotz et al. 1987). In the course of our investigation of viable alternative planning methodologies, we evaluated the applicability of two theoretical approaches: nonmonotonic logics and decision theory. In this article, we establish a simple correspondence between the two theories, describe how each theory applies to the planning task, and offer several suggestions based on the strengths and limitations of the two approaches.

Nonmonotonic Logics

One of the most compelling reasons for the use of logic as a representation language is that first-order logic does not assume how or in what context the represented knowledge will be used (Hayes 1977). Simply stated, logical inferences based on valid logical statements never result in invalid conclusions. The consistency of first-order logic is appealing, but it comes at a cost: All knowledge must be stated categorically; no possible occurrences or partially achievable goals can exist. This consistency property prevents the use of first-order logic to represent and solve problems of planning under uncertainty, which inherently involve incomplete, uncertain, and inconsistent information.

In response to this limitation, several researchers have devised augmented logical systems that allow a proposition to be assigned a truth value that signifies the proposition is consistent with the existing set of facts (true by default) even though these propositions have neither been proved nor disproved (for example, Reiter 1980). Because these default assumptions can be withdrawn based on new information, it is conceivable that new information will cause a retraction of an assertion and, therefore, a reduction in the number of provable logical statements about a particular problem. Thus, in contrast to first-order logic, the number of statements provable from a set of default assumptions does not necessarily grow monotonically with the addition of new information. Thus, these systems are called *nonmonotonic logics.*

To represent and solve a planning problem using a nonmonotonic logic, a system builder must acquire the relevant beliefs and assertions from an expert and express them in the nonmonotonic logic. The system uses a logical inference technique to construct a proof that argues for a particular course of action to satisfy the goals of the planning task. Because of their ability to express inferential uncertainty as explicit default rules, nonmonotonic logics are attractive formalisms for planning under uncertainty.

Decision Theory

Decision theory is also intended to formalize planning problems that involve substantial uncertainty (Savage 1972), but its axiomatic basis (Cox 1946; von Neumann and Morgenstern 1953) is distinct from nonmonotonic logic. In contrast to nonmonotonic logic, decision theory employs probabilities—continuous measures of belief in uncertain propositions. To accommodate conflicting problem-solving goals, quantities called *utilities* are used to denote degrees of preference toward attaining planning goals. A mathematical method for combining these probabilities and utilities is derived directly from the axiomatic basis. The solution method recommends the course of action with the maximum expected utility score. Expected utility is computed as follows for m outcomes:

$$EU(P_i) = \sum_{j=1}^{m} p(O_j | P_i) * U(O_j)$$

where $EU(P_i)$ is the expected utility of the i^{th} plan, $p(O_j|P_i)$ is the probability of the j^{th} outcome after executing the i^{th} plan, and $U(O_j)$ is the utility of the j^{th} outcome. To represent and solve a planning problem using decision theory, a system builder must acquire from an expert the probabilities, signifying the likelihood of the potential consequences of action, and the utilities, signifying the degree to which each potential consequence satisfies the problem-solving goals. The probabilities and utilities are then combined to derive an expected utility score for each plan. The plan with the highest expected utility is considered superior because its probable consequences will best satisfy the goals of the planning task. A thorough introduction to this theory of decision making is provided in von Winterfeldt and Edwards (1986) and Holloway (1979).

Limitations of the Two Formalisms

In theory, both decision theory and nonmonotonic logics can be used to solve problems of planning under uncertainty. In practice, however, both formalisms have weaknesses that restrict their use in planning systems.

Nonmonotonic Logics

By their very nature, proofs in nonmonotonic logics can later be invalidated. The temporary nature of these proofs has two important consequences for planning systems that use nonmonotonic logics. First, these systems focus on the task of constructing admissible (allowable) plans rather than selecting the optimal plan (Charniak and McDermott 1985). Second, because not all admissible plans are optimal, much of the current work on nonmonotonic logics focuses on formalizing the contexts in which nonmonotonic conclusions are appropriate, thereby precluding the construction of suboptimal plans.

Because a set of logical statements is called a *theory*, and because each consistent set of statements provable from this theory is called an *extension of the theory*, this task is sometimes called finding minimization criteria that determine the preferred extension of a default theory (Etherington 1987). These criteria are represented as logical axioms that place nonmonotonic conclusions in order of priority. This prioritization is used to determine which nonmonotonic conclusions should be considered valid in particular problem-solving settings. For example, Poole (1986) describes criteria for finding the preferred extension by giving preference to the most specific default theory. This technique is analogous to the notion that objects in a hierarchy should inherit properties from their closest ancestor.

When a nonmonotonic logic is used to represent and solve a planning problem, unintended logical inferences can result (Hanks and McDermott 1986). Several researchers (Shoham 1986; Lifschitz 1986) have developed formalizations of preference criteria that in some situations allow only intended inferences. Some recently investigated ad hoc techniques have similar goals (Cohen, et al., 1987). Doyle (1985) suggests that nonmonotonic inferences can be interpreted as a manifestation of preferences. Shoham (1987) makes this notion explicit by formally defining preference logics, which implicitly capture a partial order of intended inferences. He proposes this formalism as a generalization of other nonmonotonic logic systems and argues that the various formalizations of nonmonotonic inference are simply different methods to operationalize preferences for some inferences over others. We agree with this point and return to it later.

This view of nonmonotonic logics raises questions about their practicality. We find it counterintuitive that preferences, which are essential to the selection of intended inferences, are not explicitly represented in nonmonotonic formalisms. In addition, these formalisms offer little guidance on how such preferences should be obtained and used to influence the design of a nonmonotonic system. Finally, if new preference criteria must be devised for new problem-solving situations, this development might compromise the problem independence of the nonmonotonic logic.

Decision Theory

Whereas symbolic planning methodologies often concentrate on the construction of plans, the decision-theoretic formalism concentrates on choosing the best action from a pre-enumerated set of plans. Thus, decision theory has been used to date only as a tool to assist human planners who have already generated a small set of alternative plans to evaluate. Furthermore, decision theory has not yet been used as part of a consultation system, in part because of its inability to constrain the number of plans to evaluate and in part because of its data-hungry evaluation method. However, for problems limited in size, knowledge engineering time for decision-theoretic systems can be comparable to that for rule-based systems (Henrion and Cooley 1987). The use of even constrained decision-theoretic problem-solving models sometimes poses problems because the acquisition of probabilities and utilities is subject to systematic inaccuracies and biases (Tversky and Kahneman 1974; Nutter 1987). Although the effect of biases on problem-solving results can be estimated using sensitivity analysis (Downing 1987; Pauker and Kassirer 1980), these limitations of the knowledge-acquisition process are still a source of concern and a subject of study (Wallsten and Budescu 1983).

Another practical limitation of decision theory is the difficulty in explaining its inference procedure. The mathematical nature of maximizing expected utility does not correspond closely to human inference techniques, making this formula unappealing and unconvincing as a justification for the results of planning. A few programs have produced text explanations of quantitative problem solving (Kosy and Wise 1984; Langlotz, Shortliffe, and Fagan 1988), but this explanation process is substantially more difficult than the analogous task for symbolic reasoning programs.

In summary, we believe that decision theory is limited in its ability to constrain the set of possible plans to a tractable set of admissible plans. In contrast, nonmonotonic logics are restricted in their ability to select the optimal plan from the admissible plans. In the remainder of this article, we delineate important limitations of nonmonotonic logics with respect to the plan-selection process.

A Decision-Theoretic Model of Default Reasoning

In a previous article, we showed that decision-theoretic notions of likelihood and preference underlie rule-based planning (Langlotz, Shortliffe, and Fagan 1986). We now extend this analysis by establishing a simple correspondence between nonmonotonic logics and decision theory, thereby revealing the following characteristics of nonmonotonic planning systems. First, planning systems that employ nonmonotonic logics make implicit assumptions about problem-solving goals. Second, because goals vary among problems, nonmonotonic conclusions that are appropriate for one problem can be inappropriate for another. Third, because the appropriateness of nonmonotonic conclusions can vary, the ability of nonmonotonic logics to represent problem-independent commonsense knowledge is restricted.

We analyze nonmonotonic reasoning using decision theory to represent the decision of whether to make a nonmonotonic conclusion. Figure 1 shows how a decision tree can be used to represent decisions of this kind.[1] The tree shows two choices: to assert a nonmonotonic conclusion or not. Regardless of whether the conclusion is asserted, a probability (p) exists that the conclusion is valid, and the possibility (1 - p) remains that the conclusion is invalid.

To facilitate an intuitive understanding of the correspondence illustrated in figure 1, we first consider the example of TWEETY the bird, which is used frequently to illustrate the concept of nonmonotonic reasoning. Figure 2 shows a simple formulation of this problem in a nonmonotonic logic.[2] Initially, two propositions exist in the knowledge base. Statement 1 expresses the fact that TWEETY is known to be a bird. Statement 2 expresses if we know an object is a bird and if it is consistent with the rest of our knowledge the object can fly (that is, we cannot prove otherwise), then we are willing to add to our knowledge base the fact the object can fly. Thus, in statement 3, the conclusion is asserted that TWEETY can fly, because no specific evidence exists to the contrary.

Rich (1983) has observed that the default rule (2) is reasonable because it is highly likely that birds can fly. Our decision-theoretic representation of this problem brings up a specific question: How large a probability of error would make the nonmonotonic conclusion about TWEETY unwise? This question is addressed in detail in the following section.

Preferences and Default Reasoning

Our discussion thus far closely follows the analysis of several other researchers who acknowledge that

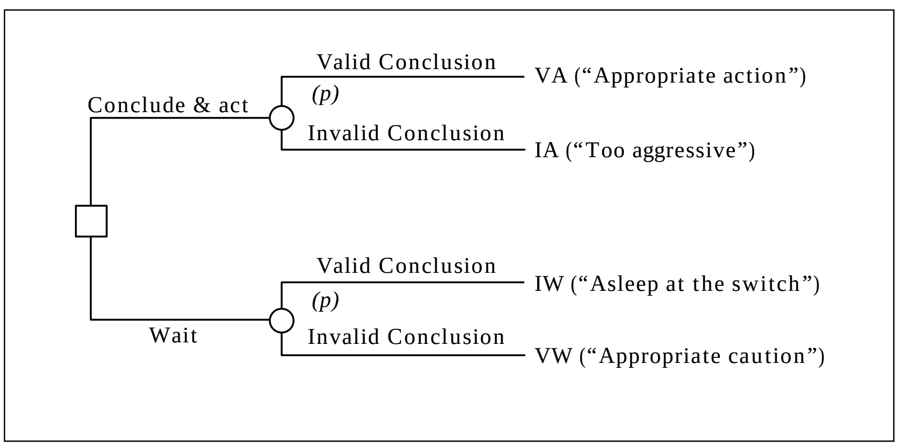

Figure 1. A Decision Tree That Represents the Decision Whether to Make a Nonmonotonic Conclusion.

The square nodes are decision nodes. The branches emanating from decision nodes represent plans among which a choice must be made. The circular nodes are chance nodes whose branches represent all the possible scenarios that might occur. The statements at the right dramatize how the problem solver might perceive each scenario in hindsight. VW signifies valid wait, IA signifies invalid action, and so on. The probability of a valid nonmonotonic conclusion is represented by p.

(1) Bird ($TWEETY$)

(2) Bird (x) and M [Fly(x)] $\supset$ Fly (x)

---

(3) Fly ($TWEETY$)

Figure 2. A Simple Proof That TWEETY Can Fly, as Expressed in Non-monotonic Logic.

M can be interpreted as "consistent with all other known facts." Therefore, M[Fly(x)] is interpreted as "the fact that x can fly is consistent with the rest of the knowledge base." The facts above the line are in the knowledge base initially, and the assertion below is proved using modus ponens.

Adapted from Reiter (1980).

nonmonotonic inference is related to notions of likelihood and probability. Rich (1983) proposes a likelihood interpretation of default logic using the certainty factor (CF) calculus (Shortliffe and Buchanan 1975).[3] Michalski and Winston (1986) proposed a similar framework for default reasoning called *variable precision logic.* The probabilistic interpretation of this default logic also corresponds to the likelihood notions we describe here.

In contrast, we concur with other researchers who propose that preference plays an important role in nonmonotonic inference. For example, Doyle (1985) identifies parallels between economic theory and nonmonotonic reasoning. He proposes that the selection of default inferences can be interpreted as a reflection of group preferences. Shoham (1987) makes a similar argument by distinguishing between first-order logics, which make conclusions because they are valid, and nonmonotonic systems, which make inferences because a reason exists to do so. He discusses several reasons that an inference might be preferred. For example, some are preferred because they represent default assumptions that facilitate communication between the system and its users. Similarly, inferences that are likely to be valid might be preferable to others; likewise, if the potential incorrectness of an inference is inconsequential, the inference might be preferable. Thus, he concludes, "One must maximize one's expected utility when selecting a nonmonotonic inference" (Shoham 1987, p. 391).

We now explore the implications of these intuitive discussions by providing a rigorous argument to support the claim that nonmonotonic inference comprises not only likelihood but also preference when considered in the context of planning problems. In planning, the task is to select or devise an appropriate action, the consequences of which are vital to the decision. Thus, implicit in the assertion of a nonmonotonic conclusion is the belief that it is acceptable to act as if the assertion were true. The desirability of acting as if an assertion were true depends not only on the likelihood of the assertion's truth but also on the desirability of the consequences of acting as if the assertion were true. (The concept of desirability is signified by the quoted statements in figure 1). Therefore, even though the likelihood that a statement is true might not change, our willingness to act as if it were true might vary. These changes in our willingness to act arise because the appropriateness of an assertion depends in part on our preferences for the potential consequences of acting on this assertion. The fact that utilities dictate the appropriateness of nonmonotonic conclusions is a recurring theme in our analysis.

This intuitive notion of utility dependence is confirmed by a mathematical analysis of the decision tree shown in figure 1. Decision theorists have derived a quantity for such trees known as the *threshold probability* (sometimes referred to as p*) (Pauker and Kassirer 1980). This quantity denotes the probability of making a valid conclusion below which a nonmonotonic conclusion is suboptimal. Similarly, the quantity 1 - p* represents the probability of error above which a nonmonotonic conclusion is suboptimal. The value of p* depends on the relative preference for each possible outcome of the problem.[4] Consider the TWEETY example posed in the context of a planning problem. How would we feel if we assumed TWEETY could fly when, in fact, TWEETY could not? We might feel comfortable with the default rule in figure 2 if we were deciding whether to buy a birdbath for TWEETY, but would we feel equally comfortable if TWEETY were a beloved pet, and we were considering dropping him off a cliff (see figure 3)? The more serious the consequences, the less comfortable we feel making the default conclusion about TWEETY. This relationship holds true even when the likelihood TWEETY can fly is invariant.

The following medical example further illustrates the implications of the relationship between the threshold probability and the utilities for the problem. Consider the following question: Does a patient with pain in the right lower abdomen suffer from appendicitis? This problem is illustrative because like the birdbath example, it is explicitly linked to action: Because the consequences of untreated appendicitis are extremely serious, all patients with known appendicitis are immediately referred for surgery. However, the surgeon cannot know with certainty before surgery whether a patient has the condition. In fact, it is considered acceptable medical practice if as high as 25 percent of patients sent to surgery do not have appendicitis (Pieper, Kager, and Nasman 1982). Thus, the threshold probability, p*, of appendicitis above which surgery should be recommended is at least 0.75. In other words, if it is evident, based on the existing laboratory and examination data, that a patient has at least a 75-percent chance of an appendicitis, the doctor should act as if the patient has appendicitis and, thus, should send the patient to surgery.

Would we feel equally comfortable assuming all birds can fly if the probability of a bird flying were only 0.75? The answer depends on what actions might be taken based on the assumption that a bird can fly. These actions depend, in turn, on the problem-solving situation. Unfortunately, a likelihood analysis of nonmonotonic reasoning would judge such a default assumption about flying birds equally acceptable to a default assumption about surgery.

Thus, the likelihood view of nonmonotonic reasoning must be extended to include the notion of utility. The general principle implied by threshold analysis is as follows: As the consequences of an invalid nonmonotonic conclusion become serious, p* increases, and the acceptable probability of error decreases. Thus, two nonmonotonic conclusions with the same probability of being invalid might not both be appropriate if the consequences of error for one are serious, and those for the other are insignificant. Consequently, we derive the following lesson from the correspondence between decision theory and default reasoning:

The appropriateness of a nonmonotonic conclusion is, in part, dependent on the (un)desirability of acting on an incorrect conclusion.

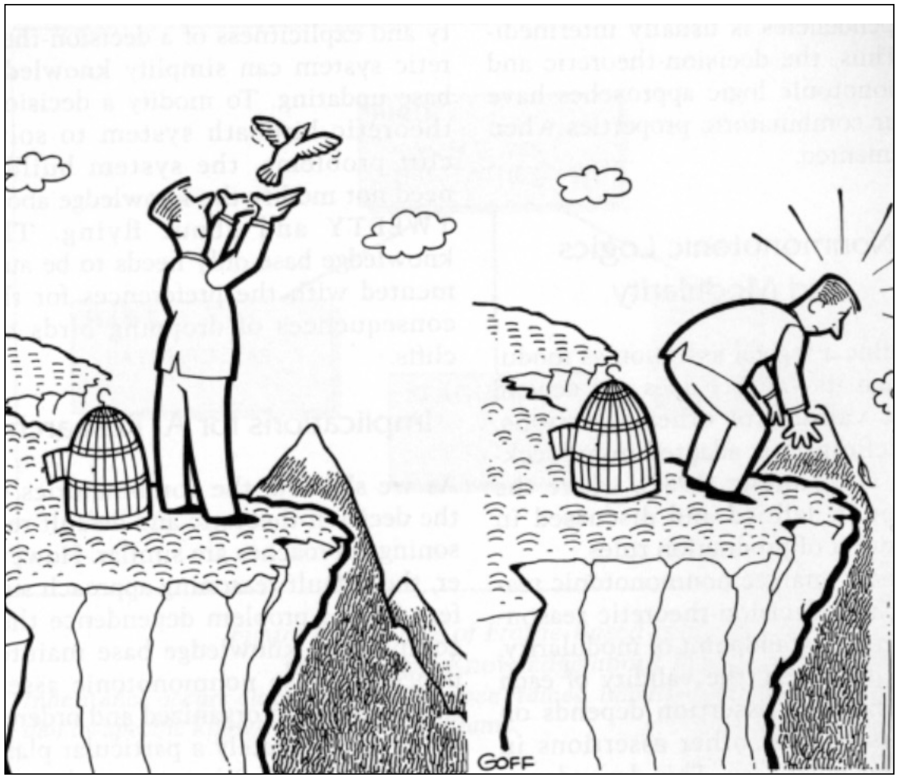

Figure 3. But, in Fact, TWEETY Can't Fly.

The undesirable consequences of such actions underlie the appropriateness of nonmonotonic conclusions.

A Practical View of the Two Approaches

We showed how a single nonmonotonic default rule, used alone, can be suboptimal. In theory, however, nonmonotonic planning systems would contain many more logical statements that, when used in concert, produce the desired behavior. We now informally analyze how these complex logical systems compare with their decision-theoretic counterparts.

Consider how the appendicitis problem might be implemented as a complete system. Additional domain knowledge would be encoded in default rules to specify those cases for which it is inappropriate to operate on a patient. For example, a nonmonotonic assertion of higher priority could be added that specified to refer for surgery when the patient has two features that suggest appendicitis: fever and high white–blood-cell count. Thus, patients would be referred to surgery unless they lacked these two indicators. This approach is oversimplified because many pieces of evidence support (and detract from) the hypothesis that the patient has appendicitis, but it illustrates the implementation strategy. A set of nonmonotonic logic statements can be constructed whose antecedents enumerate all combinations of symptoms resulting in a probability of appendicitis above the threshold, p*, and whose consequents conclude that referral to surgery is the optimal plan. To view this situation from a decision-theoretic perspective—the "gold standard" —is to conclude surgery only if the probability of appendicitis is greater than p*. As we showed, p* is computed from the undesirability of inappropriately referring a patient to surgery (or inappropriately not referring a patient to surgery).

As the number of relevant symptoms increases, the number of necessary logical statements increases dramatically. In the worst case, $O(2^n)$ logical assertions (all combinations) must be written, where n is the number of relevant pieces of evidence.[5] On average, however, a few strong pieces of evidence can dictate a decision regardless of the presence or absence of other symptoms, thereby reducing the number of logical statements required.

The computational properties of the decision-theoretic approach are similar to those of nonmonotonic logics. The decision-theoretic approach entails the use of an inference network (Pearl 1986) or influence diagram (Shachter 1986) to update the probability of appendicitis based on additional evidence. The worst case occurs when each symptom is dependent on all others: $O(2^n)$ logical statements (that is, conditional probabilities) are needed in this case. If all symptoms are independent of one another, only $O(n)$ statements are needed, each rule encoding an update in the belief in appendicitis based on the presence or absence of one symptom. In practice, the number of dependencies is usually intermediate. Thus, the decision-theoretic and nonmonotonic logic approaches have similar combinatoric properties when implemented.

Nonmonotonic Logics and Modularity

We define a logical assertion as modular when its validity does not depend on the validity of other assertions. This definition is adapted from Heckerman and Horvitz (1987), where the concept is defined and discussed in the context of production rules.

We now analyze nonmonotonic reasoning and decision-theoretic reasoning from the standpoint of modularity. As we discussed, the validity of each nonmonotonic assertion depends on the existence of other assertions in the knowledge base. This dependency is an explicitly nonmodular property of nonmonotonic logics. For example, the likelihood that TWEETY can fly is a property of our knowledge about TWEETY and rarely changes as the problem-solving goals change. In contrast, as illustrated by our discussion of the birdbath and cliff examples, the utility of assuming conclusions about TWEETY often varies among problem-solving situations (Keeney 1986). Furthermore, this knowledge is not explicitly represented. A nonmonotonic reasoning system constructed for birdbath problems would need substantial modification to ensure adequate performance on cliff problems. The interdependence (nonmodularity) among assertions makes updating and augmenting of the knowledge base far more difficult and time consuming. Thus, the performance of nonmonotonic systems is highly dependent on the problem-solving goals for which these systems were designed.

In comparison, the decision-theoretic approach maintains modularity when possible through the use of representations that explicitly encode dependencies if they exist (Pearl 1986; Shachter 1986). These representations achieve further modularity by separating likelihood knowledge from preference knowledge (that is, probabilities from utilities). The relative modularity and explicitness of a decision-theoretic system can simplify knowledge base updating. To modify a decision-theoretic birdbath system to solve cliff problems, the system builder need not modify the knowledge about TWEETY and about flying. The knowledge base only needs to be augmented with the preferences for the consequences of dropping birds off cliffs.

Implications for AI Research

As we showed, the combinatorics of the decision-theoretic and default reasoning approaches are similar. However, the default reasoning approach suffers from a problem dependence that complicates knowledge base maintenance. Because nonmonotonic assertions have been organized and ordered to reflect accurately a particular planning situation with a particular p*, the validity of these assertions is inherently dependent on the goals of the problem. Consequently, these nonmonotonic assertions tend to be specific to a particular planning problem. We now examine the implications of this conclusion for several current AI research issues: planning, commonsense knowledge, and analogic reasoning.

Use of Nonmonotonic Logics for Planning

The appeal of nonmonotonic logics is, in part, derived from the presumption that they retain the same appealing characteristic as first-order logic: The representation remains independent of the problem-solving task. As we just showed, however, the validity of a nonmonotonic assertion is inherently dependent on other assertions in the knowledge base and, thus, on the problem-specific goals. This conclusion undermines one of the frequent justifications for the use of nonmonotonic logics in the implementation of automatic problem solvers (Hayes 1977) because it shows that nonmonotonic logics have the same problem-dependent features as other representations that perform the same function.

Some researchers argue that these issues of problem dependence might be immaterial, because the effect of the (relatively problem-independent) probabilities might far outweigh the effect of the more problem-dependent utilities. We find this argument appealing. On the one hand, it advocates distinguishing and analyzing the probabilities and utilities of a problem to determine the applicability of a particular problem-solving technique. We agree with this point and view it as one of the key lessons of our analysis.[6] On the other hand, this argument implies that nonmonotonic logics are most appropriate when the utilities are relatively unimportant (that is, in low-stakes situations when the decision maker is relatively indifferent to the consequences of action). We also agree with this point; its frequent repetition would ensure that caution is exercised when nonmonotonic logic is applied to high-stakes problem-solving situations.

Commonsense Knowledge

Attempts to encode large amounts of problem-independent information (sometimes called commonsense knowledge [McCarthy 1986; Lenat, Prakash, and Shepherd 1986]) are predicated on the assumption that the knowledge representation used is problem independent (because it is isomorphic to logic). As we discussed, this property does not hold for nonmonotonic reasoning.

Consider the implications of this observation for the CYC system (Lenat, Prakash, and Shepherd 1986), which employs frame-based inheritance (a form of nonmonotonic reasoning) to encode commonsense knowledge. For example, when a knowledge engineer is encoding knowledge about birds in CYC, a bird frame might be constructed containing a can-fly slot with value true (figure 4). Several other frames, such as hawk and seagull, might inherit information from the bird frame. Other descendants of the bird frame, such as the frames for penguin and ostrich, would override the can-fly slot with the value false. When the system encounters a bird for which no more specific frame exists, it assumes that the bird can fly because general bird knowledge is assumed by default. However, there are exceptions (such as brain-damaged birds or roast turkeys). Without additional knowledge, a planning system would then act as if these invalid assumptions were true. As we noted earlier in our discussion of appendicitis, the consequences of such action can sometimes be dire.[7]

One frequent counterargument to this line of reasoning is that it ignores other information which we might have about the problem. More domain knowledge might be needed, such as frames that describe brain-damaged birds. This additional knowledge increases the number of problems for which the system's reasoning is appropriate because it lowers the probability of an unanticipated exception. The possibility for errors still exists, however (What about the Maltese falcon?). In fact, as long as we are not certain that a bird can fly, we can always imagine a problem-solving situation in which the consequences of error are so undesirable that we still do not want to make the default assumption a rule. Of course, all exceptions not explicitly represented could be excluded from our definition of bird. Then the inferences would no longer be nonmonotonic inferences; they would be logically valid.

As we saw, a knowledge engineer who encodes knowledge in a system that employs frame-based inheritance must impose an estimate of the undesirability of invalid conclusions about birds (and many other things). These judgments can represent an average of the consequences of the problem-solving situations the knowledge engineer can anticipate, or they can represent an assessment of the utilities for one particular problem. In either case, reasoning systems that might someday use the knowledge base to solve particular problems can be led astray by invalid default assumptions. Thus, the knowledge in inheritance frames in such a knowledge base cannot truly be problem independent.

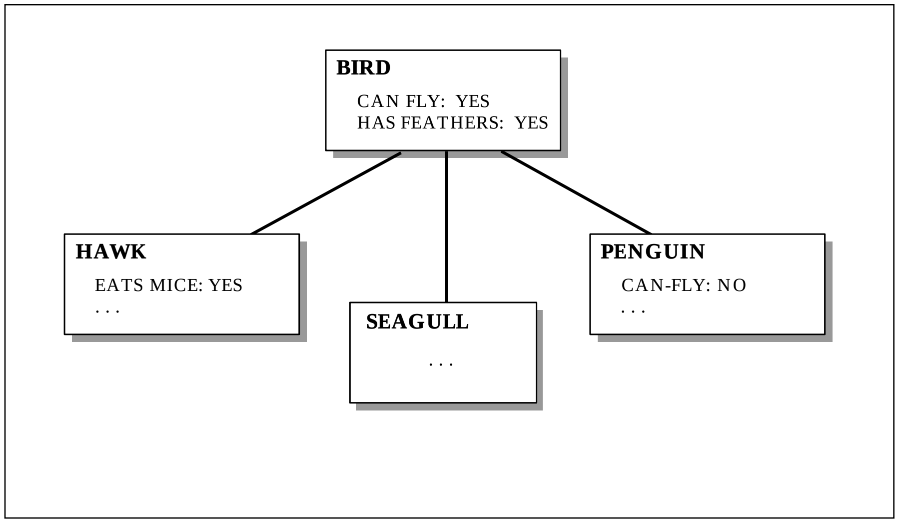

Figure 4. The Use of Frame-Based Inheritance to Represent Knowledge about Birds.

Inheritance occurs along the lines between frames. Inherited knowledge can be overridden by specific knowledge in the local frame.

Analogic Reasoning

Many researchers have proposed that the value of an analogy can be judged in terms of the likelihood that a given analogy can provide a correct solution (Russell 1986). This statement, however, doesn't consider the effect of utility information. For example, a wooden table is quite analogous to a chair if the problem solver is looking for a place to sit. However, an analogy between the table and kindling is especially valuable if the problem solver needs warmth. The aptness of these analogies might well be related to the usefulness or utility of the possible solutions as well as to their likelihood. The utility-dependent nature of analogic reasoning might be poorly understood in part because most empirical studies of analogy have examined opinions about the similarity of concepts essentially devoid of utility information, such as colors varying in lightness and saturation or circles varying in size (Shepard 1981).

Conclusions

Decision theory is designed to maintain an explicit separation between likelihood knowledge, which depends on the characteristics of the world, and utility information, which depends on the problem-solving goals. This separation provides modularity and flexibility when a knowledge base is adapted to a new planning situation and indicates that likelihood knowledge, because it is relatively utility free, might be an appropriate knowledge representation for use in commonsense knowledge bases. Thus, decision theory represents an important candidate for the implementation of planning systems.

We view our conclusions as an extension of Shoham's in that he proposes a logic formalism which identifies the preference-based nature of nonmonotonic inferences (Shoham 1987). We provide a quantitative decision-theoretic basis for this conclusion and argue that the preference-based nature of these nonmonotonic inference systems limits their utility. These limitations suggest that alternative representations might offer greater promise for the solution of problems of planning under uncertainty which require explicit representation of uncertainty and tradeoffs.

We envision decision-theoretic formalisms that can be used to guide the selection of appropriate inferences based on explicit encoding of preferences. For example, a knowledge engineer would represent commonsense knowledge in a probabilistic inference network (Pearl 1986; Shachter 1986; Lehmann 1988). Problem-specific preferences would be used to represent particular planning situations because these preferences describe the problem-solving goals and the tradeoffs among them. This utility information would be combined with the relevant likelihood-based commonsense knowledge to solve the problem at hand. Several researchers have described instances in which the separation of likelihood knowledge and utility information can be incorporated into the design of an expert system (Langlotz et al. 1987; Slagle and Hamburger 1985; Klein and Finin 1987; Langlotz 1987).

Because decision theory provides a mechanism for combining likelihood and utility knowledge during problem solving, it represents an important tool that should be considered by builders of planning systems. Used in conjunction with nonquantitative techniques to generate candidate plans (Wellman 1986), the decision-theoretic approach not only enhances the ability of a system to select the optimal plan when the explicit consideration of tradeoffs is essential but also provides a sound basis for knowledge engineering decisions. The planning problems we address indicate the continuing need to combine established normative techniques with the symbolic reasoning and representation techniques developed in AI. We believe that these combinations are fruitful and are likely to lead to enhanced decision support for problems of planning under uncertainty.

Acknowledgments

We thank Benjamin Grosof and Mark Musen for their helpful comments on earlier drafts of this article, and Lawrence Fagan for his comments on oral presentations of this work. We also thank Michael Kahn, David Klein, Michael Wellman, David Heckerman, and Eric Horvitz for many useful discussions. Support for this work was provided by the National Library of Medicine under grant no.LM-04136 and the National Science Foundation under grant no. IST83-12148. Computing facilities were provided by the SUMEX-AIM resource under National Institutes of Health grant no. RR-00785 and by the Xerox Corporation.

References

Charniak, E., and McDermott, D. 1985. *An Introduction to Artificial Intelligence*. Reading, Mass.: Addison-Wesley.

Cohen, P. R., Greenberg, M., and DeLisio, J. 1987. MU: A Development Environment for Prospective Reasoning Systems. In Proceedings of the Sixth National Conference on Artificial Intelligence, 783–788. Menlo Park, Calif.: American Association for Artificial Intelligence.

Cox, R. 1946. Probability, Frequency, and Reasonable Expectation. *American Journal of Physics* 14(1): 1–13.

Downing, K. L. 1987. Diagnostic Improvement through Qualitative Sensitivity Analysis and Aggregation. In Proceedings of the Sixth National Conference on Artificial Intelligence, 789–793. Menlo Park, Calif.: American Association for Artificial Intelligence.

Doyle, J. 1985. Reasoned Assumptions and Pareto Optimality. In Proceedings of the Ninth International Joint Conference on Artificial Intelligence, 88–90. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Etherington, D. W. 1987. Formalizing Nonmonotonic Reasoning Systems. *Artificial Intelligence* 31(1): 41–86.

Fikes, R. E., and Nilsson, N. J. 1971. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence* 2: 189–208.

Friedland, P. E., and Iwasaki, Y. 1985. The Concept and Implementation of Skeletal Plans. *Journal of Automated Reasoning* 1(2): 161–208.

Hanks, S., and McDermott, D. 1986. Default Reasoning, Non-Monotonic Logics, and the Frame Problem. In Proceedings of the Fifth National Conference on Artificial Intelligence, 328–333. Menlo Park, Calif.: American Association for Artificial Intelligence.

Hayes, P. J. 1977. In Defense of Logic. In Proceedings of the Fifth International Joint Conference on Artificial Intelligence, 559–565. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Heckerman, D. E. 1986. Probabilistic Interpretations for MYCIN's Certainty Factors. In *Uncertainty in Artificial Intelligence*, eds. J. F. Lemmer, L. N. Kanal, 167–196. New York: North Holland.

Heckerman, D. E., and Horvitz, E. J. 1987. On the Expressiveness of Rule-Based Systems for Reasoning with Uncertainty. In Proceedings of the Sixth National Conference on Artificial Intelligence, 121–126. Menlo Park, Calif.: American Association for Artificial Intelligence.

Henrion, M., and Cooley, D. R. 1987. An Experimental Comparison of Knowledge Engineering for Expert Systems and for Decision Analysis. In Proceedings of the Sixth National Conference on Artificial Intelligence, 471–476. Menlo Park, Calif.: American Association for Artificial Intelligence.

Holloway, C. A. 1979. *Decision Making under Uncertainty: Models and Choices*. Englewood Cliffs, N.J.: Prentice-Hall.

Horvitz, E. J. 1988. Reasoning under Varying and Uncertain Resource Constraints. In Proceedings of the Seventh National Conference on Artificial Intelligence, 111–116. Menlo Park, Calif.: American Association for Artificial Intelligence.

Keeney, R. L. 1986. Value-Driven Expert Systems. In Proceedings of the Conference on Multi-Attribute Decision Making via O.R.-Based Expert Systems, 141–162. Universitat Passau.

Klein, D., and Finin, T. 1987. What's in a Deep Model? A Characterization of Knowledge Depth in Intelligent Safety Systems. In Proceedings of the Tenth International Joint Conference on Artificial Intelligence, 559–562. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Kosy, D. W., and Wise, B. P. 1984. Self-Explanatory Financial Planning Models. In Proceedings of the Fourth National Conference on Artificial Intelligence, 176–181. Menlo Park, Calif.: American Association for Artificial Intelligence.

Langlotz, C. P. 1987. Axiomatic Expert Systems. In Proceedings of the Eleventh Annual Symposium on Computer Applications in Medical Care, 49–55. Washington D.C.: IEEE Computer Society Press. Also in 1989. *Computer Methods and Programs in Biomedicine*. Forthcoming.

Langlotz, C. P., Shortliffe, E. H., and Fagan, L. M. 1988. A Methodology for Generating Computer-Based Explanations of Decision-Theoretic Advice. *Medical Decision Making* 8(4): 290–303.

Langlotz, C. P., Shortliffe, E. H., and Fagan, L. M. 1986. Using Decision Theory to Justify Heuristics. In Proceedings of the Sixth National Conference on Artificial Intelligence, 215–219. Menlo Park, Calif.: American Association for Artificial Intelligence.

Langlotz, C. P., Fagan, L. M., Tu, S. W., Sikic, B. I., and Shortliffe, E. H. 1987. A Therapy Planning Architecture That Combines Decision Theory and Artificial Intelligence Techniques. *Computers and Biomedical Research* 20(3): 279–303.

Lehmann, H. P. 1988. Knowledge Acquisition for Probabilistic Expert Systems. In Proceedings of the Twelfth Annual Symposium on Computer Applications in Medical Care, 73–77. Washington D.C.: IEEE Computer Society Press.

Lenat, D., Prakash, M., and Shepherd, M. 1986. CYC: Using Common Sense Knowledge to Overcome Brittleness and Knowledge Acquisition Bottlenecks. *AI Magazine* 6(4): 65–85.

Lifschitz, V. 1986. Pointwise Circumscription: Preliminary Report. In Proceedings of the Fifth National Conference on Artificial Intelligence, 406–410. Menlo Park, Calif.: American Association for Artificial Intelligence.

McCarthy, J. 1986. Applications of Circumscription to Formalizing Common Sense Knowledge. In Proceedings of the AAAI Workshop on Nonmonotonic Reasoning, 295–324. Menlo Park, Calif.: American Association for Artificial Intelligence.

McCarthy, J. 1980. Circumscription—A Form of Non-Monotonic Reasoning. *Artificial Intelligence* 13 (1–2): 27–39.

McDermott, D., and Doyle, J. 1980. Non-Monotonic Logic I. *Artificial Intelligence* 13 (1–2): 41–72.

Michalski, R. S., and Winston, P. H. 1986. Variable Precision Logic. *Artificial Intelligence* 29(2): 121–146.

Nutter, J. T. 1987. Uncertainty and Probability. In Proceedings of the Tenth International Joint Conference on Artificial Intelligence, 373–379. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Pauker, S. G., and Kassirer, J. P. 1980. The Threshold Approach to Clinical Decision Making. *New England Journal of Medicine* 302(20): 1109–1117.

Pearl, J. 1986. Fusion, Propagation, and Structuring in Belief Networks. *Artificial Intelligence* 29(3): 241–288.

Pieper, R., Kager, L., and Nasman, P. 1982. Acute Appendicitis: A Clinical Study of 1018 Cases of Emergency Appendectomy. *Acta Chirurgica Scandinavia* 148: 51–62.

Poole, D. L. 1986. On the Comparison of Theories: Preferring the Most Specific Explanation. In Proceedings of the Fifth National Conference on Artificial Intelligence, 144–147. Menlo Park, Calif.: American Association for Artificial Intelligence.

Quinlan, J. R. 1987. Generating Production Rules from Decision Trees. In Proceedings of the Tenth International Joint Conference on Artificial Intelligence, 304–307. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Reiter, R. 1980. A Logic for Default Reasoning. *Artificial Intelligence* 13 (1–2): 81–132.

Rich, E. 1983. Default Reasoning as Likelihood Reasoning. In Proceedings of the Third National Conference on Artificial Intelligence, 348–351. Menlo Park, Calif.: American Association for Artificial Intelligence.

Russell, S. J. 1986. A Quantitative Analysis of Analogy by Similarity. In Proceedings of the Fifth National Conference on Artificial Intelligence, 284–288. Menlo Park, Calif.: American Association for Artificial Intelligence.

Savage, L. J. 1972. *The Foundations of Statistics.* New York: Dover.

Shachter, R. D. 1986. Intelligent Probabilistic Inference. In *Uncertainty in Artificial Intelligence.* New York: North Holland. Also in 1985. Proceedings of the AAAI Workshop on Uncertainty and Probability in Artificial Intelligence. Menlo Park, Calif.: American Association for Artificial Intelligence.

Shepard, R. N. 1981. Division 3 Presidential Address, American Psychological Association, Los Angeles. Also reprinted in Russell, S. J. 1986. A Quantitative Analysis of Analogy by Similarity. In Proceedings of the Fifth National Conference on Artificial Intelligence, 248–288. Menlo Park, Calif.: American Association for Artificial Intelligence.

Shoham, Y. 1987. Nonmonotonic Logics: Meaning and Utility. In Proceedings of the Tenth International Joint Conference on Artificial Intelligence, 388–393. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Shoham, Y. 1986. Chronological Ignorance: Time, Nonmonotonicity, Necessity, and Causal Theories. In Proceedings of the Fifth National Conference on Artificial Intelligence, 389–393. Menlo Park, Calif.: American Association for Artificial Intelligence.

Shortliffe, E. H., and Buchanan, B. G. 1975. A Model of Inexact Reasoning in Medicine. *Mathematical Biosciences* 23: 351–379.

Slagle, J. R., and Hamburger, H. 1985. An Expert System for a Resource Allocation Problem. *Communications of the ACM* 28(9): 994–1004.

Tversky, A., and Kahneman, D. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science* 185: 1124-1131.

von Neumann, J., and Morgenstern, O. 1953. *Theory of Games and Economic Behavior.* New York: Wiley.

von Winterfeldt, D., and Edwards, W. 1986. *Decision Analysis and Behavioral Research.* New York: Cambridge University Press.

Wallsten, T. S., and Budescu, D. V. 1983. Encoding Subjective Probabilities: A Psychological and Psychometric Review. *Management Science* 29(2): 151–173.

Wellman, M. 1986. Qualitative Probabilistic Networks for Planning under Uncertainty. In Proceedings of the AAAI Workshop on Uncertainty in Artificial Intelligence, 311–318. Menlo Park, Calif.: American Association for Artificial Intelligence.

Notes

1. The term decision tree has sometimes been used to denote a recursive structure that represents a sequential classification process (Quinlan 1987). We use the term in its traditional sense—to represent the possible scenarios that arise when we plan under uncertainty.

2.We illustrate our approach using a variant of Reiter's default logic (Reiter 1980). Our basic argument also applies to other formalisms for default reasoning, including circumscription (McCarthy 1980), McDermott and Doyle's nonmonotonic modal logic (McDermott and Doyle 1980), and several other formalisms.

3. Heckerman (1986) has subsequently related the CF calculus to probability theory.

4. The derivation of p* is based on the observation that at the threshold, the decision maker is indifferent between the two choices. Consequently, Pauker and Kassirer derive p* by mathematically expressing the expected utility of each choice in terms of the probabilities and utilities involved, then setting these expressions equal to one another and solving for p.

5. For the purposes of this combinatoric analysis, we assume that evidence is constrained to be either present or absent.

6. A formal framework for addressing these issues was proposed by Horvitz (1988).

7. Shoham (1987) makes a similar point when he discusses of the chances of being stabbed while walking on the street.

Curtis P. Langlotz is a National Library of Medicine predoctoral fellow in Medical Informatics at Stanford University. In June of this year, he will complete an M.D. and a Ph.D. in Medical Information Science and begin residency training in diagnostic radiology. His research interests include planning under uncertainty, automated explanation, and the integration of decision theory and AI techniques.

Edward H. Shortliffe is associate professor of Medicine and Computer Science at Stanford University. He is the scientific director of the Medical Computer Science Group in the Knowledge Systems Laboratory and chief of the Division of General Internal Medicine at Stanford University School of Medicine. His research interests include models for evidential reasoning and representation techniques to support advanced explanation capabilities.