Reasoning about Beliefs and Actions under Computational Resource Constraints[1]

Eric J. Horvitz

Medical Computer Science Group
Knowledge Systems Laboratory
Stanford University
Stanford, California 94305-5479

Although many investigators affirm a desire to build reasoning systems that behave consistently with the axiomatic basis defined by probability theory and utility theory, limited resources for engineering and computation can make a complete normative analysis impossible. We attempt to move discussion beyond the debate over the scope of problems that can be handled effectively to cases where it is clear that there are insufficient computational resources to perform an analysis deemed as complete. Under these conditions, we stress the importance of considering the expected costs and benefits of applying alternative approximation procedures and heuristics for computation and knowledge acquisition. We discuss how knowledge about the structure of user utility can be used to control value tradeoffs for tailoring inference to alternative contexts. We address the notion of real-time rationality, focusing on the application of knowledge about the expected timewise-refinement abilities of reasoning strategies to balance the benefits of additional computation with the costs of acting with a partial result. We discuss the benefits of applying decision theory to control the solution of difficult problems given limitations and uncertainty in reasoning resources.

1 Introduction

Enthusiasm about the use of computation for decision support and automated control within high-stakes domains such as medicine has stimulated interest in the construction of systems that behave consistently with a coherent theory of rational beliefs and actions. Numerous investigators interested in the automation of uncertain reasoning have converged on the theoretical adequacy of the decision-theoretic basis for rational action [11, 3, 17]. Recent discussions about computational approaches to reasoning with uncertainty have centered on the degree to which probability and utility theory can handle inference problems of realistic complexity. Investigators have answered criticism about the inadequate expressiveness of probability theory by pointing out that the normative basis addresses consistent inference with measures of belief and preference, not issues surrounding the formulation of problems [18]. Other researchers have shown that probability theory and utility theory are logically equivalent to the satisfaction of a small set of intuitive properties [35, 18]. Still others have responded to complaints of intractability by demonstrating techniques that can solve

[1]This article appeared originally in the *Proceedings of the Third Workshop on Uncertainty in Artificial Intelligence*, Seattle WA, July 1987. AAAI and Association for Uncertainty in Artificial Intelligence, Mountain View, CA, pp. 429-444. This work was supported by Grant NCC 2-220-51 from the NASA-Ames Research Center and Grant RO1LM04529 from the National Library of Medicine. Computer facilities were provided by the SUMEX-AIM Resource under Grant RR-00785 from the National Institutes of Health.

relatively complex real-world problems [29, 13].

In this paper, we move beyond discussions of the degree to which the theories of probability and utility are able to solve real-world problems. We focus on situations where it is clear that insufficient resources prohibit the use of the normative basis for a complete analysis. That is, we are interested in studying cases where normative reasoning is clearly inadequate because of pressing resource limitations. We are concerned with rational strategies for handling such resource insufficiencies. We have been exploring resource constraints and the feasibility of normative approaches to knowledge assessment, computation, and explanation under scarce resources [15, 14]. That is, we have sought to apply the principles of normative rationality to reason *about* the solution of a base-level decision problem.

We focus our attention on real-time decision making. Resource-constraint issues can be especially salient in the context of real-time requirements. In the real world, delaying an action is often costly. Thus, computation about belief and action often incurs inference-related costs. The time required by a reasoning system for inference varies depending on the complexity of the problem at hand. Likewise, the costs associated with delayed action vary depending on the stakes and urgency of the decision context. The real-time problem is further complicated by the existence of uncertainty in the cost functions associated with delayed action. We are searching for uncertain-reasoning strategies that can respond flexibly to variations in the availability of resources. The intent of our research is to develop coherent approaches to generating and selecting the most promising strategy for particular problem-solving challenges.

Notions of *bounded rationality* that have been exercised in earlier discussion of intelligent systems faced with complex problems shun a formal perspective as too costly [31, 23]. Most research on reasoning and acting under resource constraints has focused on the discovery of relatively simple satisficing approaches to problem solving. These approaches may stray far from the levels of utility that might be achieved through the pursuit of more sophisticated normative analyses. Losses may be especially significant in high-stakes decision making. Rather than reject the pursuit of a theoretical foundation for ideal belief and action, we seek to extend coherently the principles of normative rationality to situations of uncertain, varying, and scarce reasoning resources. Although the research path may be fraught with challenging theoretical problems, potential benefits of the work include the construction of artifacts that have greater expected value and the elucidation of principles of reasoning under scarce resources.

2 Components of Uncertain Reasoning

We have found it useful to decompose uncertain reasoning into three components: problem formulation, belief entailment, and decision making. *Problem formulation* is the task of modeling or constructing the reasoning problem. This task often involves enumeration of relevant hypotheses and dependencies among hypotheses. There are no formal theories for problem formulation; in many machine-intelligence projects, engineers charge domain experts with the task of enumerating all relevant propositions and of structuring the dependencies among the propositions. *Belief entailment* or *inference* is the process of updating measures of truth assigned to alternative hypotheses as new evidence is uncovered. In most schemes, the degree of truth or *belief* in the presence of a hypothesis can range continuously between complete truth and complete falsity. Such belief-entailment schemes include probability theory [18, 27], fuzzy logic [38], Dempster–Shafer theory [30], and MYCIN certainty factors [2]. Finally, *decision making* is the process of selecting the best action to take. A decision or action is an irrevocable allocation of valuable resources.

The classical decision-theoretic basis defines rational beliefs and actions with the axioms of probability theory and utility theory. Probability theory dictates that the assignment and entailment of belief in the truth of propositions should be consistent with a parsimonious set of axioms. The logical equivalence of these axioms with a small set of intuitive properties desired in a measure of belief has been demonstrated [7, 18]. *Utility theory* [36] dictates the consistent assignment and updating of the *value* of alternative actions given the values of alternative outcomes and the degrees of belief in the outcomes. Measures of value consistent with the axioms of utility theory are called *utilities*. Von Neumann and Morgenstern, the authors of utility theory, proved that agents making decisions consistent with the axioms of utility would behave as though they associate utility values with alternative outcomes and would act to maximize their expected utility [36].

The application of probability theory for belief assignment and utility theory for decision making defines a *normative basis* for reasoning under uncertainty. The term *normative* refers to the notion that probability theory and utility theory have been accepted in several disciplines as a consistent axiomatic basis for inference that is considered optimal. That is, for many people, the normative framework defines a *rational* theory for belief and action.

3 The Limited Scope of the Normative Basis

Artificial-intelligence research has highlighted the problems that lurk beyond the axiomatic framework defined by probability and utility theory. The real-world problems examined by machine-intelligence investigators are often more complex than were the problems previously tackled with decision theory. In applying the normative basis to problems of real-world complexity, the limited domain of discourse of the theory becomes apparent. It is clear that significant aspects of problem modeling and inference in the real world are absent from the language and axioms of the normative basis. The normative theory's sole focus on the consistent assignment of measures of belief and preference is dwarfed by the complex task of constructing and solving the uncertainty problem. For example, the axioms have nothing to say about the modeling process. They do not address issues surrounding the most appropriate propositions to represent, the level of abstraction to select, or the degree of completeness or detail of interdependencies to represent.

The normative basis also does not address the most appropriate inference technique for reasoning under specified constraints in computational resources. Classical applications of normative rationality have implicitly assumed sufficient computational resources for reasoning about an optimal action; the basis itself does not explicitly address issues surrounding the value of alternative approaches to generating partial solutions in reasoning systems that might be dominated by varying limitations in computational or engineering resources.

There is much research to be done on reformulating problems and inference strategies deemed optimal in a world with infinite resources for performance in resource-limited environments. In this regard, we see promise in the development of techniques for examining alternative models and inference strategies as the *objects* of design- and real-time normative metalevel analysis. This task involves determining, in a tractable fashion, the most promising expenditure of engineering or computational resources. Our research has highlighted the notion that a system with the ability to reason under uncertainty about complex real-world problems often requires extensive knowledge about the domain at hand as well as knowledge about the expected behavior of alternative inference strategies.

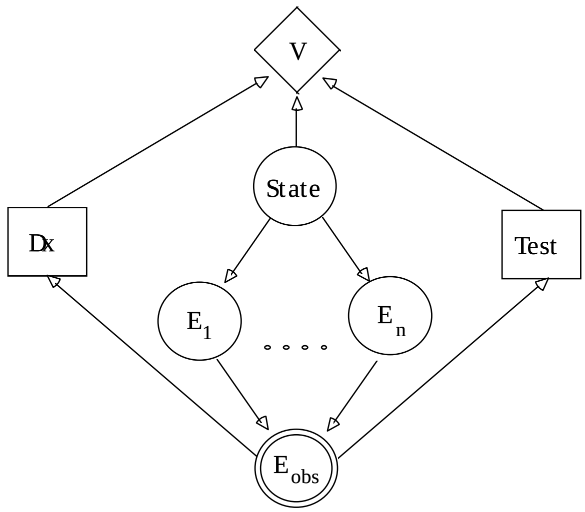

Figure 1: A decision network representing a step in the sequential diagnosis problem.

4 The Complexity of Rational Inference

Let us pause briefly to consider the complexity of normative rationality. Recent research has focused on the computational complexity of probabilistic reasoning. The research has been based on analyses of uncertain-reasoning problems represented with graphs. The most popular representation uses directed graphs to represent explicitly conditional dependencies and independencies among beliefs in propositions [28, 5, 27]. Many researchers have ascribed a common semantics to the directed graphs. The representation is often called a *belief network*. In a belief network, an arc between a node representing proposition A and another node representing proposition B expresses knowledge that the probability distribution over the values of B depends on the specific values of proposition A. If there is no arc from A to B, the probability distribution for B is not directly dependent on the values of A. Less expressive representations commonly employed in artificial-intelligence research have not allowed specific independencies to be represented efficiently [12].[1]

Belief networks are special cases of more general graphical representations that can represent available *actions* and the *utility* of alternative outcomes in addition to beliefs [21, 29]. These graphs have been called *influence diagrams* and *decision networks*. An example of a portion of a simple decision problem for medical diagnosis is shown in Figure 1. The node labeled E_{obs} represents the current state of the observed evidence. The evidence in this case consists of symptoms that are *caused* by the pathophysiologic state, represented by the node labeled $State$. The decisions are represented by the square nodes labeled Dx and $Test$. The utility of the state of affairs to a patient is represented by the diamond-shaped value node, labeled V. In this case, the utility depends on the actual disease present ($State$), the decision about additional testing ($Test$), and the decision to assume a particular diagnosis (Dx). In the real world, there are frequently dependencies among the observations and disease state, making this a complex inference problem.

Although the directed-graph representations allow the expression of inference problems that can be solved efficiently, many topologies have resisted tractable algorithmic solution. A troubling topology is the *multiply connected network* [27]. Such inference problems belong to a class of difficult problems that have been shown to be $\mathcal{NP}$-hard [6]. Problems in complex areas such as medicine often require representation with multiply connected networks. Thus, in the worst case, rational beliefs and actions demand computation that is exponential in the size of the problem.

It is clear that many uncertain-reasoning problems require more computation time than may be available before a commitment to action is required. What can be done when the cost of inference becomes intolerable? As a first step, investigators might search for special-case inference techniques designed for the efficient solution of specific problem types (e.g., specific belief-network topologies or belief distributions). However, proofs, such as the one demonstrating the worst-case intractability of multiply-connected networks, diminish hope that special methods will be discovered for solving important classes of problems. For many situations, we will need to develop intelligent approximation procedures and heuristics that focus the expenditure of resources on the most relevant aspects of the uncertain reasoning problem at hand.[2]

The pressures of complex decision making in real-time force Bayesian theoreticians and engineers to consider alternatives to normative reasoning: Under time constraints (or other resource constraints, such as the cost of knowledge-acquisition), approximations and more poorly characterized heuristic techniques often have a higher expected value than does complete normative reasoning. The delay associated with inference might be so costly that an approximation method or heuristic procedure might have a greater expected value, in spite of assured suboptimality or uncertainty in the performance of the strategy. Thus, constraints in resource can transform a non-normative technique into the "preferred choice" of devout Bayesians, and can convert the strictest formalists into admirers of heuristics.

We have been investigating the problem of reasoning under specified constraints within the Protos project. A focus of this research addresses the use of decision theory for selecting among alternative problem-formulation and inference strategies. We believe that the representation of explicit knowledge about the costs associated with computation, such as time delay, will be useful in complex uncertain reasoning problems. Although we hope to discover approximate inference techniques that show clear dominance, we believe that it is often important to reason about inference tradeoffs under uncertainty at the metalevel.

5 Inference under Resource Constraints

Simple normative reasoning systems have been constructed based on a single model constructed as a static basis and acted on by a single inference strategy. We are interested in techniques for reformulating a base problem into one that will be of greater value than a complete analysis would be, given computational resource constraints. A reasoning system with knowledge about the behavior of alternative approximation methods and heuristics, and about the costs associated with inference-based delay, might provide valuable computation under resource constraints. A complete normative analysis of the same problem might be a worthless or costly enterprise.

We shall now raise several issues about strategies that can focus computational attention on the most relevant portions of uncertain reasoning problems. Challenging components of this research include the development of approximation procedures and heuristics that are insensitive to small variations in resource availability, the representation of knowledge about the value structure of the problem, and the development of compiled and real-time control strategies that can recognize problems, understand the problem-solving context, and select or construct the most valuable inference strategy.

5.1 Integration of Knowledge about Inference-Related Costs

Theoretical models of rationality must include the costs associated with rational inference itself. We wish to include knowledge about the reasoner in the reasoning problem. The representation of inference costs can be valuable in the control of inference. A crucial aspect of integrating knowledge about real-world costs, benefits, and tradeoffs into a reasoning system is the acquisition of knowledge about the value of important attributes of computer performance to the users of computer systems.

We have found it useful to decompose the value associated with computational inference into two components. We assert that the application of an inference strategy is associated with some net benefit or cost to an agent–such as a system user, a robot, or a computational subsystem–that relies on computation for decision making. We use the term *comprehensive value of computation* V_c to refer to the *expected utility* associated with the application of a computational strategy. This value is a function of the strategy, of the problem at hand, of the best default action in response to the problem, and of the problem-solving context.

We view the comprehensive value as having two components: the *object-related* value and *inference-related* value.[3] The *object-related* value V_o is the expected utility associated with the best action or result available to an agent, given a state of the world. In the computational setting, it is often useful to reason about marginal increases in object-level value with computation. For example, changes in the object-related value associated with the use of an expert system for assistance with a complex medical diagnosis problem refers to the costs and benefits associated with the change in information about the entities in the medical problem such as treatment alternatives, likelihoods of possible outcomes, and costs of recommended tests. The *inference-related* value V_i is the expected disutility intrinsically associated with *computation*, such as the cost a physician might attribute to the delay of a decision because of the time required by an expert system to generate a recommendation, or the cost associated with his inability to understand the rationale behind a machine's recommendation. The *net value of computation*, ΔV_c, refers to the change in the comprehensive value, given some quantity of computation. This is just the difference between the increase in object-level utility and the cost of the additional computation.

Knowledge about costs and benefits of computation can be integrated into the decision-network representation; the influence-diagram representation is expressive enough to capture the base and metalevel control problem. A more comprehensive representation of our simple diagnosis problem is portrayed in Figure 2. We have added an influence diagram capturing the uncertainty and possible decisions regarding the control of reasoning in the object-level decision problem represented by the original decision network. The metalevel control problem can make use of partial knowledge about the performance of alternative strategies to choose the best procedure for solving the base diagnosis problem. The arcs and nodes of the control problem represent knowledge about base-level computation and autoepistemic knowledge about the costs of metareasoning. The traditional value node from the base problem corresponds to the object-level utility V_o in the control network. The metalevel reasoning problem is to optimize the comprehensive value V_c through considering the object-level value in addition to the inference-related factors of time delay and deadlines. The node *Clarity* indicates that the *transparency* and explainability of inference may be another important dimension of inference-related cost in the context of an expert system [19, 16].

5.2 Multiple Attributes of Inference and Outcome

As indicated in Figure 2, the costs of reasoning may have several components. Multiple dimensions of utility can be ascertained through consideration of attributes of outcomes and problem solving that are important to a computational agent or system user. Such multiattribute utility can be assessed through analysis of the values of possible scenarios. Frequently, a function can be constructed that captures the relationships among attributes of computational value in important contexts. The value assigned to alternative computational behaviors often can be described by a qualitative or more detailed function that represents the relationships among important dimensions of the perceived costs and benefits associated with alternative outcomes. Such value functions assign

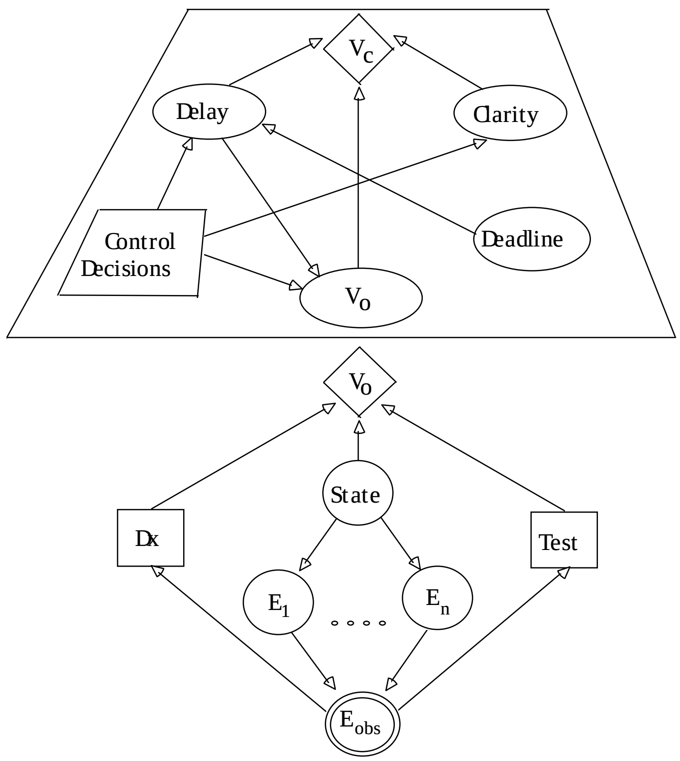

Figure 2: The addition of a decision network that represents costs, benefits, and decisions about alternative approaches to solving the object-level diagnosis problem. The base problem and the control network together represent a richer decision problem that includes knowledge about belief and preferences regarding the problem-solving process itself.

a single utility measure to computation based on the status of a vector of attributes.

As an example, a planner attempting to maximize a robot's expected utility in a complex environment will generally have to consider multiple components of value in computational goals. In making decisions about its next set of goals, a robot may have to consider the distance and accessibility of the essential staples of electrical power and oil, the positions of alternative crate-stacking tasks, the speed with which it can create a new plan, and its distance from other robots that might require or lend assistance. Similarly, the value associated with the use of a medical expert system in a particular context might be a function of a number of attributes, including speed of computation, accuracy of recommendation, and clarity of explanation. We have been working with expert physicians in the intensive-care and tissue-pathology domains to ascertain value models relating measures of utility to multiple attributes of computation. There is almost always uncertainty in the object-level state that results from the expenditure of resources. Thus, in the general case, we must sum over a probability distribution of object-level attributes to generate an expected comprehensive utility. In terms of net comprehensive utility,

$$\Delta V_c = \sum_{\vec{v}\,'} V_c(\vec{v}\,', \vec{r}\,')\; p(\vec{v}\,'|\vec{r}\,') \;-\; V_c(\vec{v}, \vec{r})$$

where $\vec{v}$ and $\vec{r}$, respectively, are the vectors of object-level and inference-related attributes without additional computation, and $\vec{v}'$ and $\vec{r}'$ are the revised vectors, associated with additional computation.

5.3 Inference Tradeoffs

Computation in a world of bounded resources often is associated with cost/benefit *tradeoffs*. With a computational tradeoff, the benefit associated with an increase in the quantity of one or more desired attributes of computational value is intrinsically linked to costs incurred through changes imposed on other attributes. More specifically, we define a tradeoff as a relationship among two attributes of utility, such as the *immediacy* and *precision* of a computational result, each having a positive influence on the perceived total value of computer performance, and each constrained to be a monotonically decreasing function of the other over some relevant range. In the case of our sample tradeoff,

$$PRECISION = \mathcal{F}(IMMEDIACY), \;\; t_o \leq IMMEDIACY \leq t_n$$

where $\mathcal{F}$ is some monotonically decreasing function over the range bounded by computational time delays t_o and t_n. This definition can be generalized to the case where the value assigned to tuples of a subset of relevant attributes is a monotonically decreasing function of tuples composed of other attributes. The tradeoff between the immediacy and the precision or accuracy of a solution is particularly explicit in methods that incrementally refine a computational result with time. Most reasoning systems have been designed with implicit assumptions about the handling of inference tradeoffs. The expected value of an automated reasoner's behavior, within dynamic environments, might be greatly enhanced by endowing the system with the ability to tailor inference to a range of problems and contexts.

5.4 Earlier Research

Protos research is pursuing the decision-theoretic control of computational problem solving. Work is focused on the task of the rational control of decision-theoretic inference and of several sorting and searching tasks. Related machine-intelligence research on the control of reasoning by Smith, and by Treitel and Genesereth, has explored the usefulness of applying utility theory in the selection of alternative logical reasoning strategies [32, 34]. Our work differs from the logic theorem-proving work in its pursuit of theoretical tools for the control of computation under scarce resources, of decision-theoretic inference under resource constraints, and of the structure of partial results and multiple dimensions of utility.

Previous research in decision science has touched on the formal integration of the costs of reasoning into decision-making inference. The earliest discussion of the explicit integration of the costs of inference within the framework of normative rationality was introduced by Good [9], who made the distinction between what he referred to as *type I* and *type II* rationality. Good defined type I rationality as inference that is consistent with the axioms of decision theory without regard to the cost of inference. Type II rationality is behavior that takes into consideration the costs of reasoning. We carry out our work in the spirit of Good; that is, we seek a rational approach to rational inference under scarce resources. Other related work in decision science centered on the value of analysis. These studies have explored the likely benefit of doing a decision analysis or continuing to refine a decision. Matheson [24] explored the value of spending additional effort to analyze a bidding problem proposed by Howard [20]. Watson and Brown [37] and Lindley [22] describe issues surrounding the application of a preliminary decision analysis to assist in the decision as to whether an individual should embark on a more costly decision analysis of the base problem at hand.

6 Toward a Timewise-Refinement Paradigm

Classical approaches to normative inference have pursued the determination of point probabilities. In fact, the complexity proof described in Section 4 is based on the assumption that point probabilities are required. The classical interest in calculating final answers permeates computer science. Complexity theorists have focused almost exclusively on proving results about the time and space resources that must be expended to run algorithms to termination [8, 1, 25]. In the real world, strict limitations and variations on the time available for problem solving suggest that the focus on time complexity for algorithmic termination is limited; analyses centering on how *good* a solution can be found in the time available for computation are of importance.

The major rationale for dwelling on the time complexity of algorithmic termination resides in the simplifying notion in algorithms research that a computer-generated result can be assigned only one of two measures of utility: either a solution is found and is of value, or a solution is not found and the result is therefore valueless. However, it is often possible to enumerate representations and inference techniques that can provide partial solutions that have varying *degrees* of value. For example, we could consider the value of different types of partially sorted files instead of dwelling on our inability to always complete a sort under uncertain and varying time limitations [14].

An approach to developing techniques for optimizing the value of uncertain reasoning under ranging resource limitations is the development of problem reformulation and inference schemes that allow the generation and efficient manipulation of *partial results*. We are interested in representation and reasoning methods that allow a result to be refined with increasing amounts of computation. In analyzing the timewise refinement behavior of algorithms, it is crucial to consider knowledge about the *value structure* of partial results. We believe that formalizing the costs and benefits, and the cost–benefit tradeoffs, associated with inference in differing contexts will be beneficial in the development of insights about useful approximations and heuristics.

6.1 Describing Resource Limitations

Before we discuss several properties of inference that are desirable for problem solving under bounded resources, let us focus more closely on reasoning resources and resource availability. A *resource* is some costly commodity required for inference, such as computation time or memory. In addition, information that can be used to direct problem solving, or an object-level action, is a reasoning resource. Decisions about resource expenditure span the dynamic allocation of time and memory at runtime, as well as the costs, benefits, and amortization of design-time investments in knowledge acquisition and hardware capabilities. We can consider any modifiable dimension of hardware or problem-solving architecture as a resource that might be extended with some economic expenditure. For example, a processor clock rate or an instruction set is a commodity that might be enhanced. The discussion in this paper focuses on issues surrounding the dynamic allocation of scarce resources, given a predefined reasoning system. However, in the general case, we must consider how longterm investments in hardware capabilities affect optimal resource allocation strategies in the short term.

Although reasoning resource is generally a multiattribute vector, we focus here on computation time. We define the resource required by a computational agent to solve a problem completely, using a reasoning strategy within the context of an agent's capabilities, as the *complete resources* $r_c(S, P, \xi)$, where S is a reasoning strategy, P is a problem instance defined by a decision model and context, and ξ is the background information about an agent's composition, capabilities, and state of information.[4] ξ includes an agent's hardware makeup, problem-solving architecture, and knowledge about problem solving. We define the object-related value associated with the complete solution of a problem in a context C as the *ideal object-related value*, $V_o^*(P, C)$. We term the resource applied to solving a problem the *allocated resources*, r_a. We refer to the ratio of the allocated and complete resources as the *resource fraction*, $r_f(S, P, \xi)$. The resource fraction can be a useful metric for reasoning about computation under bounded resources. We can use the notions of resource fraction, comprehensive value, object-related value, and inference-related value to express properties desired of inference within environments dominated by varying resource limitations.

6.2 Desiderata of Bounded-Resource Computation

What properties of inference are desirable for doing inference under uncertain challenges and scarce computational resources? We seek representation and control strategies that configure knowledge and processing in a manner that is effective in light of uncertainty in the amount of resources available for computation. For example, we desire representation and inference methodologies that allow the most relevant updating to occur early on. Also, as many real-world applications may involve reasoning under variations in the amount of time available for inference, it is desirable to design inference strategies that are insensitive to small ranges in resource fraction. We will now focus on two classes of desiderata that are useful in reasoning under scarce resources.

6.2.1 Flexibility

The first desiderata address the desirability of a graceful response to diminishing resource levels. These properties may be summarized with the concept of flexibility. We use *flexibility* in the context of reasoning under scarce resources to refer to the ability to react gracefully to a range of problems and resource availabilities. Flexible inference strategies are able to generate a wide variety of custom-tailored responses. Flexible responses are an important dimension of intelligence. They are crucial in planning for action under uncertainty. Flexible inference can be especially useful in light of uncertain resource constraints. Uncertainty about problem solving plagues simple agents immersed in complex, competitive environments; constraints on an agent's reasoning and representation resources lead to uncertainties about the problems that may be faced and about the performance of alternative reasoning strategies in solving the problems. Flexible strategies allow limited agents to benefit greatly by reacting robustly to challenges that may be only *partially characterized* ahead of time. Several desired properties of flexible computation are as follows:

• *Value Continuity* We desire the comprehensive value of computation, the object-related utility, and the inference-related utility of a strategy to be continuous functions of the resource fraction as that fraction ranges from zero to one. That is,

$$\lim_{r_f \to r_f'} V_c(S, P, C, r_f, \xi) = V_c(S, P, C, r_f', \xi)$$

where $r_f' > r_f$ and $V_c(S, P, C, r_f, \xi)$ is the comprehensive value of computation associated with an agent's applying inference strategy S to problem P in context C with resource fraction r_f[5].

• *Value Monotonicity* We desire the object-related utility of a strategy to be a monotonically increasing function of the resource fraction as that fraction ranges from zero to one. We refer to the continuous decrease of object-related value with decreasing allocation of resource over ranges that show a net positive value of computation as *graceful degradation*.

• *Convergence* We desire strategies that demonstrate convergence on the optimal object-related value at some level of resource expenditure, $r_c(S, P, \xi)$,

$$\lim_{r \to r_c} V_o(S, P, C, r_f, \xi) = V_o^*(P, C)$$

Strategies that show continuity and monotonicity are typically less efficient for the complete solution of a problem than a discontinuous strategy because of the overhead associated with the generation of valuable partial results.

• *Value Dominance* We seek problem-solving strategies with value-dominant intervals over available quantities of resource within real-world problem-solving contexts. We define *value-dominant* intervals as ranges of resource fraction over which the gain in the comprehensive value of computation is a monotonically increasing function of resource.

6.2.2 Bounded Optimality

Another class of desiderata for systems reasoning under scarce resources addresses a desire for rational control of reasoning and action. We use the term *bounded optimality* to refer to the optimization of computational utility given a set of assumptions about expected problems and constraints in reasoning resources. We can construct different classes of bounded optimality and pursue approximations to these definitions. We define alternative types of bounded optimality by making explicit assumptions about the nature of utility and about the composition and information state, ξ, of a computational agent. For example, we may wish to pursue behavior that is bounded optimal over some duration, considering the frequencies of different problems. In this case, we seek to optimize an agent's utility for some time frame in the context of a distribution of problem challenges.

We note that bounded optimality may not be well defined for arbitrary problem-solving contexts. Problems with the definition of bounded optimality include the difficulty of probing belief without changing it,[6] and the possible sensitivity of control decisions to inference at an arbitrary level of metareasoning. For local optimization of utility in real-time, the discovery and confirmation of bounded optimality could add substantial computation costs, thereby decreasing the value of the computed result. Thus, in the context of short-term optimization of utility within complex environments, verification of simple notions of bounded optimality will generally rely on partial analyses at design-time.

The determination of true bounded optimality requires proving lower-bounds on the solution of problems given the informational and computational constraints at hand. In the absence of theoretical limits, we can reason about the relative bounded optimality of agents limited to a distinct set of reasoning strategies. This perspective is useful, given current research on the solution of probabilistic-inference problems with alternative approximation strategies.

• *Bounded Strategic Optimality* We desire a reasoning system to apply strategies from its repertory of strategies such that its expected utility is a maximum, given probability distributions over the costs and benefits of applying alternative strategies. A tuple of strategies S should be selected such that the agent's comprehensive value is maximized. That is,

$$S^* = \arg\max_S [\max_r V_c(S, P, C, r, \xi)]$$

Strategies available to an agent include that of ceasing computation and taking physical action. A system seeking to satisfy bounded strategic optimality captures notions of limited rationality under resource constraints in terms of a specific problem instance. Such a reasoner would attempt to optimize the comprehensive value of its computation and physical activity, regardless of the method lying at the foundations of its inference. We could modify the definition of strategic optimality by adding additional constraints. For example, we might impose a bound on the proportion of reasoning resources an agent could apply to real-time metalevel reasoning.

We can extend the local nature of bounded strategic optimality by considering the expected utility associated with solving a distribution of problems, expected over a period of time. Such a perspective can be useful in comparing the effectiveness of agents, with different compositions and abilities, immersed in distinct problem contexts. For example, given a set of agents, an environment C, and time horizon t, we may prefer the behavior of an agent, A^*, where

$$A^*(C, t) = \arg\max_A \sum_{i=1}^{n} t \ f_{P_i}(C) * [V_d(A, P_i, C, \xi) + \Delta V_c(S^*[A, P_i, C, \xi])]$$

where $f_{P_i}(C)$ is frequency of problem type P_i in context C; V_d is the expected value associated with an agent taking it's best default action in response to a problem challenge; ΔV_c is the incremental value of the best strategy available to the agent. This *independent-challenge model*, and other definitions of agency preference, can be useful in reasoning about such factors as the relative performance of different agents in specific contexts, the value of learning in a domain, and the utility of adding a new capability to an agent's problem-solving repertory. In the independent-challenge model, we assume independence among problems, and consider the distribution of problems as independent of the type of agent, and of the agent's abilities to solve problems. The description also assumes that the utility assigned to the performance of an agent in solving a challenge is independent of the time the challenge is posed. More detailed analyses include the representation of dependencies among actions, problems, and utilities, and a consideration of the manner in which expenditures and actions, made at different times, are valuated. A comprehensive analysis of the net value of an agent should additionally include terms capturing expenditures for the initial configuration of an agent's hardware, for ongoing hardware maintenance, and for knowledge acquisition. Finally, in reasoning about the absolute utility gains derived from the use of a computational agent, it is important to consider the decision-making policies that would be undertaken in the absence of the agent.

7 Flexible Inference and Intelligent Control

Two promising areas of research on rationality under resource constraints are (1) the development and characterization of intrinsically flexible inference strategies, and (2) the mastery of techniques and computational architectures for efficient decision-theoretic control. The puruit of innovation in both areas will highlight principles of bounded-optimal problem solving, and greatly facilitate the development of robust autonomous agents and decision-support systems that are oriented to human preferences.

7.1 Promising Probabilistic Inference Approaches

Several classes of approximation methods and heuristics are promising sources of useful strategies for bounded-resource computation.

• *Bound Calculation and Propagation* There has been ongoing interest in the calculation of upper and lower *bounds* on point probabilities of interest [5]. Probabilistic bounding techniques determine bounds on probabilities through a logical analysis of constraints acquired from a partial analysis. Bounds become tighter as additional constraints are brought into consideration. Cooper has applied a best-first search algorithm to focus attention on the most relevant aspects of the problem in calculating bounds on the hypotheses [5].

• *Stochastic Simulation* Simulation techniques are approximation strategies that report a probability distribution or partial characterization of a distribution over probabilities of interest through a process of weighted random sampling [13, 26]. In many cases, the distribution over the probabilities is approximated by the *binomial* distribution. The variance with which the distribution converges on a probability with additional computation depends on the topology of the network, and on the nature of the probabilistic dependencies within the network. Recent work has shown current simulation algorithms to have intolerably slow convergence rates in many realistic cases [4] . Stochastic simulation is nevertheless a promising class of inference for the derivation of useful bounded-resource computation strategies.

• *Completeness Modulation* Completeness-modulation strategies center on techniques for reasoning about attributes of the uncertain-reasoning model to include in an analysis. Completeness modulation can be used to simplify the topology of a belief network through deletion of classes of dependencies. In one form of completeness modulation, arcs in the graph are assigned priorities by heuristic measures of context-dependent "importance" that capture the benefits of including the dependencies in alternative contexts. Such heuristic measures may be encoded during knowledge acquisition. The measures allow a reasoning system dynamically to construct a model that will be subjected to some inference procedure (e.g. bounding, simulation, complete normative analysis). Under time constraints, a completeness-modulation approach can allow components of the problem viewed as most important to be included in an analysis early-on. We have worked with experts to acquire measures of *importance* on probabilistic dependency a medical domain.[7] A long-standing heuristic in reasoning under uncertainty involves the default assumption (or the imposition) of conditional independence among propositions considered by a system. Specific dependencies are included in an overwhelming independent model when they become salient. Such an assumption greatly reduces the resources required for knowledge assessment and computation. Assumptions of global independence have been made in many reasoning systems that have been deemed to perform adequately (e.g., the MYCIN certainty-factor model [11, 17] and the early probabilistic diagnostic programs [10, 33]). The actual costs and benefits of defaulting to conditional independence among evidence in many real-world problems have not been determined. A promising area of research is the prioritization of classes of dependencies and their inclusion in a model as computational resource becomes available.

• *Abstraction Modulation* In many cases, it may be more useful to do normative inference on a model that is deemed to be complete at a particular level of abstraction than it is to do an approximate or heuristic analysis of a model that is too large to be analyzed under specific resource constraints. It may prove useful in many cases to store several belief-network representations, each containing propositions at different levels of abstraction. In many domains, models at higher levels of abstraction are more tractable. As the time available for computation decreases, network modules of increasing abstraction can be employed.

• *Local Reformulation* Local reformulation is the modification of specific troublesome topologies in a belief network. Approximation methods and heuristics designed to modify the microstructure of belief networks will undoubtedly be useful in the tractable solution of large uncertain-reasoning problems. Such strategies might be best applied at knowledge-encoding time. An example of a potentially useful local reformulation is the use of tractable *prototypical dependency* structures, such as the noisy-OR structure [27]. The benefits of using such structures for knowledge acquisition and inference could warrant the use of tractable prototypical dependencies in situations where the latter are clearly only an approximation of more complex dependencies.

• *Default Reasoning and Compilation* Under severe time pressure, general default beliefs and policies may have more expected value than does a computed result. Indeed, in some application areas, it may be useful to focus a reasoning system's scope of expectation through the compilation and efficient indexing of computed advice for actions of great importance or of high frequency, or that are needed in time-critical situations. The idle-time solution of problems, directed by likelihood, importance, and criticality, can generate libraries of compiled belief and action. The relative worth of storing heuristic default knowledge or precomputed beliefs and policies depends on several factors, including the tractability of available inference strategies, the nature of the available resource fraction, and the complexity of expected outcomes in the application area. Decisions about whether to compute or to store recommendations may also be quite sensitive to the specific costs of computer memory and knowledge assessment. Careful consideration of the value structure of components of computation in real time and in system-engineering settings can help to elucidate specific cases of such tradeoffs.

7.2 Decision-Theoretic Control

The different categories of inference described in the preceding section highlight the notion that a system may be able to choose among alternative strategies and strategy sequences to generate useful bounded-resource solution strategies. The efficient solution of complex problems under ranging resource constraints will require management of the costs, benefits, and uncertainties associated with applying portions of available reasoning resources to object and metalevel reasoning.[8] Useful conceptions of bounded optimality can be built with decision theory for decisions about reasoning and computational resources in design-time, idle-time, and real-time settings. Beyond its use for directing probabilistic inference, decision-theoretic control promises to be useful for optimizing the value of a broad variety of computational tasks, such as sorting and searching. As control decision making makes use of uncertain knowledge about problem solving, a number of interesting research issues arise that relate to the acquisition and use of partial characterizations of performance. There is much opportunity for developing automated techniques for learning and applying relevant information about problem-solving performance. Other promising research focuses on the development of methods for recognizing problems and for monitoring problem-solving progress.

We have found it useful in our work to simplify the task of controlling reasoning under bounded resources by assuming different prototypical classes of resource constraints within an application area. Multiple inference approaches and representations of a problem can be devised, each tailored for the long-term maximization of V_c in contexts with specific distributions of resources and challenges. We are currently studying the value of metalevel reasoners with access to several base-level strategies and with rich control knowledge about the value of the strategies in different problem contexts. It appears useful to construct computational architectures that grant a control reasoner easy access to separate knowledge bases, including a strategy base that contains information about

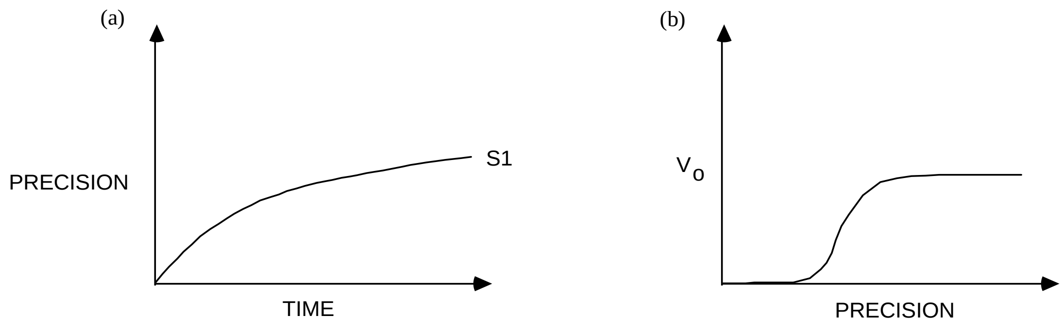

Figure 3: (a) The timewise refinement of the precision of a result with additional computation. (b) The object-level utility of the refined result.

problem-solving performance along relevant dimensions of computation, a utility base that holds information about the construction of utility through combination of the dimensions of performance, and a resource base that contains information about the costs of delay in prototypical resource contexts.

8 Reasoning about the Time-Precision Tradeoff

We now demonstrate several of the concepts we have presented with an example. We focus on the use of knowledge about multiple components of value at the metalevel to tailor inference to the appropriate context. The example reflects ongoing work on inference under bounded resources [14]. Although the results can be derived formally, we shall describe the sample problem with a set of qualitative curves for clarity. The curves capture important functional relationships among components of computational value in alternative contexts.

Consider an inference problem from one of our application areas: An automated control system is faced with a rapidly evolving set of respiratory symptoms in a patient in an intensive-care unit. Assume that our system's action depends on $p(C|E)$—the probability of a condition C given the observed symptoms E. In particular, this probability is important in deciding whether or not the system will respond by advising that a patient be given a costly or dangerous treatment for condition C.

What kinds of strategies might our autonomous pulmonary decision-making system employ to respond rationally under pressing time constraints? Assume that the system has a base model deemed during the construction of a system by a human expert to be an *adequately complete model* of the relevant world. Figure 3(a) demonstrates the knowledge that the medical decision system may have about the expected rate of computational refinement of the *precision* of the requested probability for a probabilistic bounding strategy, $S1$, given this type of problem. A measure of precision could be the variance of the second-order distribution over the probability of interest. This probability of interest is the number that would be generated given infinite computation. In the case of a probabilistic bounding scheme we might interpret the bounds as a uniform distribution over the final probability.

Let us now introduce utility considerations. The assignment of value to computed partial results of increased precision depends on the decision context; the value of an imprecise probability to a user can range greatly, depending on the end use of the probabilistic information. A system could be endowed with knowledge about the changes in expected value of perfect information with

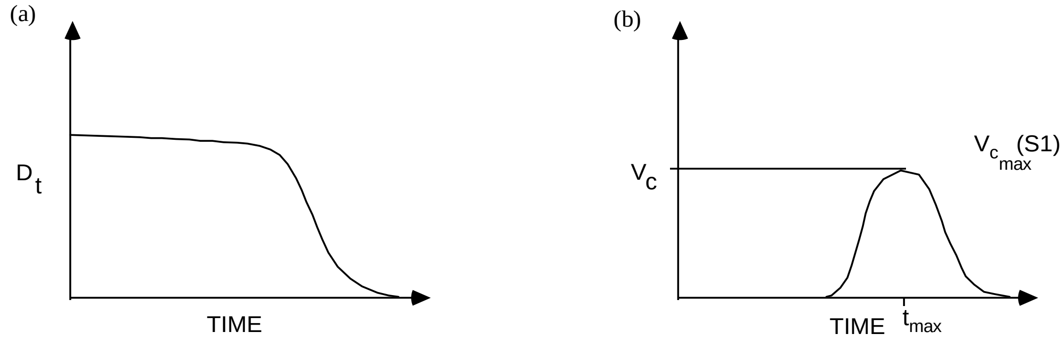

Figure 4: (a) The inference-related cost based in delay of action. (b) The comprehensive value of computation. The maximum value of computation is reached at time t_{max}.

additional inference. To encode knowledge about the assignment of object-related value to partial probabilistic results of different precisions, we could work with an expert to assess the utility directly, apply some preenumerated value function, or formally analyze the decision-making context. Let us briefly examine the last option.

Utility theory dictates that the object-related value (V_o) in this binary decision problem is determined by the probabilities and utilities of four possible outcomes: the patient either has or does not have the condition, and the system will either treat or not treat for the condition. Simple algebraic manipulation can be used to show that the optimal object-related value of information depends on the costs associated with treating a person without the condition, the benefits of treating a person with the condition, and the probability of the condition being present. Thus, changes in the distribution over the actual probability of the pathophysiologic state can be assigned a measure of value within the decision-theoretic framework.

Assume that our expert system has actively acquired information about the context in which the desired probability will be used, and has characterized the object-related value of the probability of the condition as a function of the precision of the reported probability. A plausible value function for this situation is shown in Figure 3(b). The function demonstrates that the rate of refinement of the object-related value can vary greatly with increasing precision.

So far, we have examined only object-related value considerations. In the real world, time delay can be quite costly. While we have been dwelling on issues surrounding the refinement of the object-related value, our patient has been gasping for breath. In this case, it is clear that, for any fixed measure of object-related value, the comprehensive value of the result decreases with the amount of time that a user must wait for that result to become available. It is thus important for a medical decision system to have knowledge of the inference-related utility associated with computational inference.

Let us assume that a physician with extensive knowledge about the realm of possibility in the intensive-care unit had, at an earlier date, represented context-specific knowledge about the rate at which the object-related value should be discounted with the passage of time. That is, utility assessment at the time of knowledge engineering revealed that the expert physician's preferences about the cost of delay in such a context could be represented as an independent multiplicative discounting factor, D_t, ranging in value between one and zero with the passage of time. We have considered this factor independent for simplicity of presentation; such a discount rate may depend on the status of the probabilities and outcomes. In this example, we have framed inference-related knowledge acquisition at the level of *classes of criticality* associated with unresolved pathophysiol-

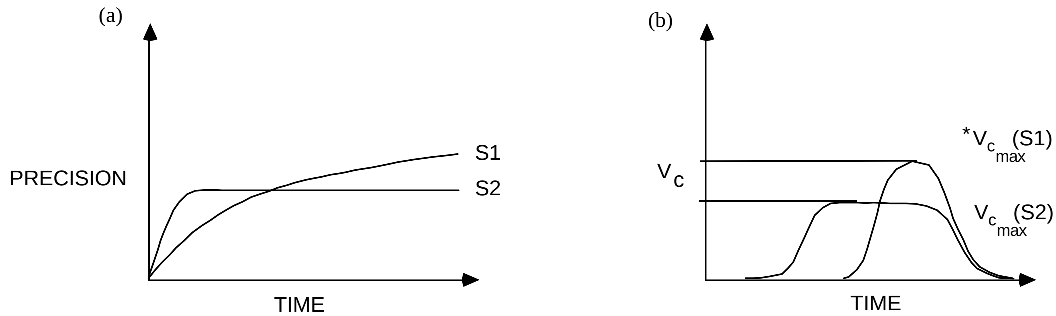

Figure 5: (a) The timewise refinement of another inference strategy. (b) A comparision of the comprehensive value of computation of the two inference strategies.

ogy. Thus, the object-related value is multiplied by the inference-related utility-discounting factor to generate the net value of an answer as time passes. A function demonstrating such a degradation of the utility of the analysis with time is shown in Figure 4(a).

If the information in the three functions is combined, the comprehensive value, V_c, of the computational process to a system user as a function of time can be derived. This result is displayed in Figure 4(b). Notice that the comprehensive value has a global maximum $V_{c_{max}}$ at a particular time, t_{max}. This is the period of time the computer system should apply inference scheme $S1$ to maximize the value to the patient of its reasoning. Although spending additional time on the problem will further increase the precision, the comprehensive value to the user will begin to decrease. Integrating a consideration of the cost and benefits of computation into an analysis of probabilistic inference makes it clear that the cost of computation can render the solution of the complete problem inappropriate.

8.1 Metalevel Reasoning about Alternative Strategies

So far, we have considered characteristics of the computational value of only one reasoning strategy. Assume that the system's metalevel reasoner has knowledge about the existence of another inference strategy, $S2$, based on stochastic simulation. Assume further that the expected precision over time of the second strategy is represented by the curve portrayed in Figure 5(a). Finally, assume that the system has knowledge that, within this context, the strategy has a higher expected rate of refinement of precision early on, but a lower long-range rate of refinement than that of stochastic simulation.

If we apply the same object-value and inference-related functions presented previously to the new inference strategy, we can derive a new comprehensive value function. This function is shown in comparison to the previously derived comprehensive value function in Figure 5(b). Notice that $V_{c_{max}}(S1) > V_{c_{max}}(S2)$. A control strategy satisfying the *bounded strategic optimality* property would select strategy $S1$ given all current knowledge about available probabilistic-inference strategies and about the decision context at hand.

8.2 Contraction of the Decision Horizon

Now, suppose that the decision context has changed in a way that affects only the inference-related cost function describing the discounting of object-related value with time. In the new context, we have a much sharper discounting of the object-related value with time, as shown in

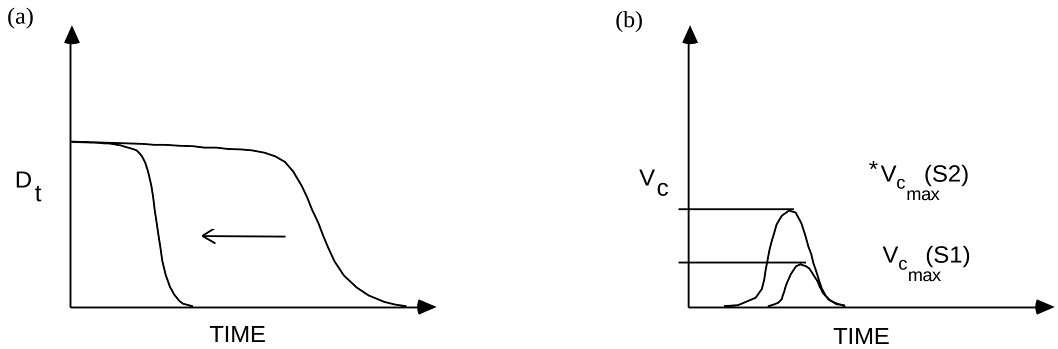

Figure 6: (a) A contraction of the decision horizon. (b) A new dominance in a more time-critical context.

Figure 6(a). Such a decreased decision horizon may be associated with situations requiring rapid response, as might be the case when our patient suddenly begins to show signs of critically low oxygenation. If we derive the comprehensive value functions for inference strategies $S1$ and $S2$ with the new object-related value-discounting function, we see a new dominance. Figure 6(b) shows that $V_{c_{max}}(S2) > V_{c_{max}}(S1)$. That is, in contexts of greater time criticality, the value achieved by strategy $S2$ will dominate that achieved by $S1$, and thus $S2$ will now be the strategy of choice.

8.3 Defaulting to Default Actions

We have dwelled on strategies that can provide partial results through computation. Before concluding, we shall move beyond uncertain inference to examine default reasoning. The default strategy $S3$ is shown in Figure 7(a). As portrayed in the figure, a default rule for a particular context often can be made available with relatively little computation. For example, a set of default-action rules could be efficiently stored and indexed by propositions that determine the relevance of each rule.

Notice, in Figure 7(a), that the object-related value of the default strategy within a problem context does not change with time; after being made available, the object-related value of a default strategy is not refined with computation. In this case, we portray the maximum object-related value of the default rule that would "fire" in the context at hand as being a fraction of the object-related value attainable through computation.

A compiled policy with low object-related value could be the strategy of choice in situations of extreme time criticality. For example, if our patient's blood pressure were suddenly to fall greatly, a theoretically suboptimal "compiled" default strategy requiring little computation might dominate. We depict graphs reflecting this situation in Figure 7(c) and (d).

We have described the simple example of diagnosis under conditions of pressing time constraints to demonstrate how a reasoning system can apply knowledge about the value of alternative strategies to optimize the value of computation to a system user. The example demonstrates how normative reasoning techniques might be applied to control reasoning to select the best strategy for solving an inference problem under different resource constraints.

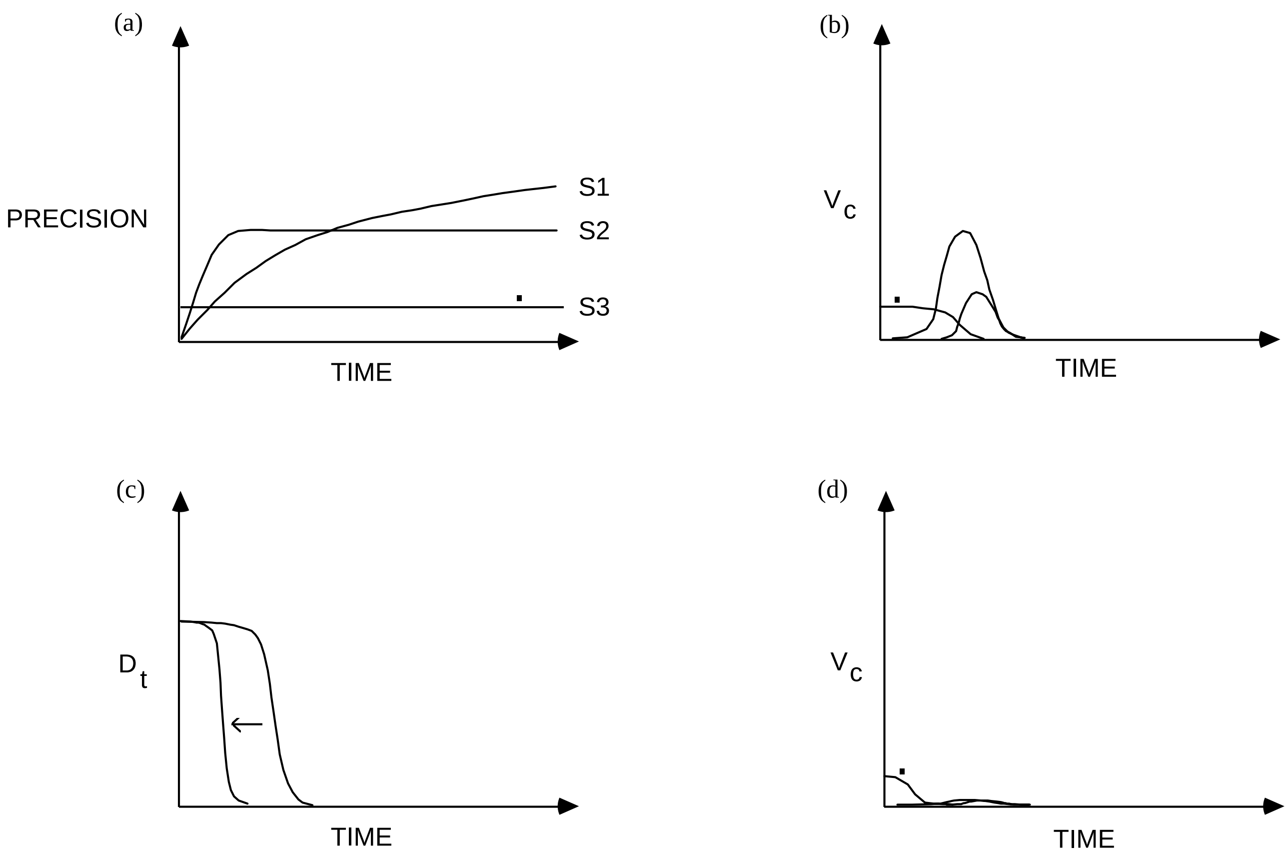

Figure 7: (a) A default reasoning strategy (marked). (b) The comprehensive value of the default strategy. (c) Another shift in the decision horizon. (d) The dominance of the utility of the default strategy in a more critical decision-making context.

9 Summary

We reviewed several issues surrounding a normative approach to beliefs and actions under resource constraints. We began the paper with a discussion of the limited scope of the normative basis for reasoning under uncertainty in the real world. We described the application of knowledge about inference-related costs in systems that reason under uncertainty, touching on the assignment of measures of utility to multiple attributes of computation and the notion of computational tradeoffs. After discussing properties capturing flexibility and optimality desired in bounded-resource inference, we presented classes of approximation procedures and heuristics that promise to be useful in reasoning under resource constraints. Finally, we presented an example that is representative of continuing investigation of the costs and benefits of alternative inference strategies in different settings. There is great opportunity for applying decision theory to make rational control decisions based on uncertain knowledge about multiple dimensions of problem-solving performance. Associated research problems include the design-time, idle-time, and real-time identification of features of problems that have relevance for problem-solving decisions; the automated discovery and representation of problem-solving performance; and the development of useful classes of bounded-optimal problem solving. We believe that continuing research on principles of representation and control of reasoning under conditions of varying computational and engineering resources will be crucial for building systems that can act effectively in the real world.

Acknowledgments

John Breese, Greg Cooper, Bruce Buchanan, David Heckerman, Ronald Howard, Nils Nilsson, Edward Shortliffe, and Patrick Suppes have provided useful feedback on this research.

Notes

[1]For example, the popular rule-based representation may encourage researchers to make global assumptions about the absence of dependencies among propositions.

[2]We use *approximation* to refer to strategies that generate results with a categorical margin of error; we use *heuristic* to refer to strategies with ill-characterized or uncertain performance. A strategy may be viewed as heuristic in terms of specific dimensions of its behavior. According to the perspective of a heuristic as a strategy characterized by a probability distribution over its performance on a set of problem instances, investigation, leading to new bounds on the behavior of a strategy, can transform a heuristic method into an approximation strategy.

[3]In general, we may have to consider dependencies between the object- and inference-related value. We assume a function f that relates V_c to V_o, V_i, and additional information about the problem-specific dependencies that may exist between the two components of value—that is, $V_c(\alpha, \beta, \gamma) = f[V_o(\alpha, \gamma), V_i(\beta, \gamma)]$, where α and β represent parameters that influence the object- and inference-related utilities respectively, and γ represents the parameters that influence both the object- and the inference-related utilities.

[4]All representations can be viewed as incomplete to some extent; we use *complete* to refer to an object-related model perceived by a system designer to be an adequate representation of a problem. Although a set of techniques for pruning models, referred to as *sensitivity analysis*, has been elucidated, no formal techniques for verifying the completeness of models have been developed.

[5] *Value continuity* may be generalized to *bounded discontinuity*, where some upper bound on an ϵ

change in V_c is specified for some δ change in r_f over ranges of resource fraction. The statement of such constraints or of a probability distribution over such constraints can be used as a partial characterization of heuristic strategies for important aspects of performance.

[6]The problem with the process of measurement or reflection affecting the state of an agent's belief, analogous to Heisenberg's uncertainty principle regarding the problem of measurement interfering with the state of the measured physical phenomenon, arose in discussions with David Heckerman.

[7]An importance metric may also be useful in directing the allocation of resources during learning and knowledge assessment.

[8]We seek theoretically sound means of determining the extent and level of metareasoning. However, in many cases, empirically determined or heuristically assumed limits on metalevel effort will have to be imposed; if not, an agent may be faced with the prospect of infinite analytical regress.

References

[1] A.V. Aho, J.E. Hopcroft, and J.D. Ullman. *Data Structures and Algorithms.* Addison-Wesley, Menlo Park, California, 1983.

[2] B.G. Buchanan and E.H. Shortliffe, editors. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project.* Addison-Wesley, Reading, MA, 1984.

[3] P. Cheeseman. In defense of probability. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence, Los Angeles, CA*, page ?? Internation Joint Conference on Artificial Intelligence, ?? 1985.

[4] H.L. Chin and G.F. Cooper. Bayesian belief network inference using simulation. In L.N. Kanal, J.F. Lemmer, and T.S. Levitt, editors, *Uncertainty in Artificial Intelligence 3*, pages 129–148. North Holland, New York, 1989.

[5] G.F. Cooper. *NESTOR: A Computer-based Medical Diagnostic Aid that Integrates Causal and Probabilistic Knowledge.* PhD thesis, Computer Science Department, Stanford University, Stanford, CA, November 1984. Rep. No. STAN-CS-84-48. Also numbered HPP-84-48.

[6] G.F. Cooper. Probabilistic inference using belief networks is NP-hard. *Artificial Intelligence*, 42:393–405, 1990.

[7] R. Cox. Probability, frequency and reasonable expectation. *American Journal of Physics*, 14:1–13, 1946.

[8] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W.H. Freeman and Company, New York, 1979.

[9] I.J. Good. Rational decisions. *J. R. Statist. Soc. B*, 14:107–114, 1952.

[10] G.A. Gorry. Computer-assisted clinical decision making. *Methods of Information in Medicine*, 12:45–51, 1973.

[11] D.E. Heckerman. Probabilistic interpretations for MYCIN's certainty factors. In L.N. Kanal and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence*, pages 167–196. North Holland, New York, 1986.

[12] D.E. Heckerman and E.J. Horvitz. On the expressiveness of rule-based systems for reasoning under uncertainty. In *Proceedings AAAI-87 Sixth National Conference on Artificial Intelligence, Seattle, WA*, pages 121–126. Morgan Kaufmann, San Mateo, CA, July 1987.

[13] M. Henrion. Propagation of uncertainty by probabilistic logic sampling in Bayes' networks. In J.F. Lemmer and L.N. Kanal, editors, *Uncertainty in Artificial Intelligence 2*, pages 149–164. North Holland, New York, 1988.

[14] E.J. Horvitz. Reasoning about inference tradeoffs in a world of bounded resources. Technical Report KSL-86-55, Medical Computer Science Group, Section on Medical Informatics, Stanford University, Stanford, CA, March 1986.

[15] E.J. Horvitz. Toward a science of expert systems. *Proceedings of the 18th Symposium on the Interface of Computer Science and Statistics*, pages 45–52, March 1986.

[16] E.J. Horvitz. A multiattribute utility approach to inference understandability and explanation. Technical Report KSL-28-87, Medical Computer Science Group, Section on Medical Informatics, Stanford University, Stanford, CA, March 1987.

[17] E.J. Horvitz and D.E. Heckerman. The inconsistent use of measures of certainty in artificial intelligence research. In L.N. Kanal and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence*, pages 137–151. North Holland, New York, 1986.

[18] E.J. Horvitz, D.E. Heckerman, and C.P. Langlotz. A framework for comparing alternative formalisms for plausible reasoning. In *Proceedings AAAI-86 Fifth National Conference on Artificial Intelligence, Philadelphia, PA*, pages 210–214. Morgan Kaufmann, San Mateo, CA, August 1986.

[19] E.J. Horvitz, D.E. Heckerman, B.N. Nathwani, and L.M. Fagan. The use of a heuristic problem-solving hierarchy to facilitate the explanation of hypothesis-directed reasoning. In *Proceedings of Medinfo, Washington, DC*, pages 27–31. North Holland, New York, October 1986.

[20] R.A. Howard. Value of information lotteries. *IEEE Transactions of Systems Science and Cybernetics*, SSC-3(1):54–60, 1967.

[21] R.A. Howard and J.E. Matheson, editors. *Readings on the Principles and Applications of Decision Analysis*. Strategic Decisions Group, Menlo Park, Ca., 1984.

[22] D.V. Lindley. Reconciliation of decision analyses. *Operations Research*, 34:289–295, 1986.

[23] J.G. March. Bounded rationality, ambiguity, and the engineering of choice. *Bell Journal of Economics*, pages 587–608, 1978.

[24] J.E. Matheson. The value of analysis and computation. *IEEE Transactions on Systems Science, and Cybernetics*, 4:211–219, 1968.

[25] C.H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1982.

[26] J. Pearl. Evidential reasoning using stochastic simulation of causal models. Technical Report R-68, CSD-8600, Cognitive Systems Laboratory, UCLA Computer Science Department, September 1986.

[27] J. Pearl. Fusion, propagation, and structuring in belief networks. *Artificial Intelligence*, 29:241–288, 1986.

[28] W.F. Rousseau. A method for computing probabilities in complex situations. Technical Report 6252-2, Center for Systems Research, Stanford University, Stanford, CA, May 1968.

[29] R.D. Shachter. Evaluating influence diagrams. *Operations Research*, 34:871–882, 1986.

[30] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, NJ, 1976.

[31] H.A. Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, 69:99–118, 1955.

[32] D.E. Smith. Controlling inference. Technical Report STAN-CS-86-1107, Computer Science Department, Stanford University, April 1986.

[33] P. Szolovits and S.G. Pauker. Categorical and probabilistic reasoning in medical diagnosis. *Artificial Intelligence*, 11:115–144, 1978.

[34] R. Treitel and M.R. Genesereth. Choosing directions for rules. In *Proceedings AAAI-86 Fifth National Conference on Artificial Intelligence, Philadelphia, PA*, pages 153–157. Morgan Kaufmann, San Mateo, CA, August 1986.

[35] M. Tribus. *Rational Descriptions, Decisions, and Designs.* Pergamon Press, New York, 1969.

[36] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior.* Princeton University Press, Princeton, NJ, 1947.

[37] S.R. Watson and R.V. Brown. The valuation of decision analysis. *J.R. Statist. Soc. A.*, 141(1):69–78, 1978.

[38] L.A. Zadeh. The role of fuzzy logic in the management of uncertainty in expert systems. *Fuzzy Sets and Systems*, 11:199–227, 1983.