# Some Fundamental Problems and Opportunities from the Standpoint of Rational Agency*

**Eric J. Horvitz**
Medical Computer Science Group
Knowledge Systems Laboratory
Stanford University
Stanford, California 94305

March 3, 1989

## 1 Introduction

Growing enthusiasm about the application of autonomous reasoning in high-stakes domains like medicine and aerospace has stimulated interest in systems that behave in accordance with a coherent theory of rationality. Within such domains, the losses associated with suboptimal decisions tend to render simple satisficing approaches inadequate and provide incentive for attempting to optimize the utility of computational activity. In pursuing research on rational agency over the last several years, a number of problems have come to be highlighted as rich areas for future research. In this paper, we review promising prospects for future study. The research topics will be motivated by recent research and preliminary theoretical and empirical results.

At the heart of intelligent behavior is the pursuit of maximal utility by reasoners with relatively limited representational and inferential abilities. Constraints on an agent's reasoning and representation resources lead to inescapable uncertainties about the problems that may be faced and about the performance of alternative reasoning strategies in solving the problems.

*This work was supported by a NASA Fellowship under Grant NCC-220-51, by the National Science Foundation under Grant IRI-8703710, and by the National Library of Medicine under Grant RO1LM0429. Computing facilities were provided by the SUMEX-AIM Resource under NIH Grant RR-00785.

Almost any interesting work on intelligent problem solving centers on the development of techniques that simple artifacts immersed in complex environments can depend on for sucessful competition with environment challenges and with threats by other intelligent competitors for precious commodities. Uncertainty and runs deep at several levels of intelligent foundations in the complexity of the problems. However, an intelligent reasoner must turn inward and reason about the deep uncertainties about the structure of utility, about the representations that it must construct, and about the nature of solving inference problems that are posed by the models.

As background, one focus of activity in our research lab has been the exploration of decision-theoretic reasoning strategies in medicine. In fact our initial foray into research on rationality under resource constraints was motivated by problems with the effective application of decision theory in expert systems and in intelligent monitors for medicine. Current research focuses on the control of decision-theoretic inference itself with decision-theory. However, our research has also dwelled on the decision-theoretic control of other types of problems, and on foundational problems with computational approaches to rationality.

# 2 Bounded Rationality to Bounded Optimality

## 2.1 Intractability and Heuristics

Our work has been captured by the pursuit of bounded optimality. Decision theory has not been a popular tool in machine intelligence research. Notions of *bounded rationality* that have been exercised in many discussion of intelligent systems faced with complex problems shun a formal perspective as too costly[9,7]. The work of Simon, March, and others, in the 1950s pointed out problems with the complexity of using probabilities and utilities for handling difficult problems. Simon's was interpreted by many, as a charge to seek out and study ill-characterized heuristic strategies. And, in fact, that interpretation still dominates much of AI.

## 2.2 Toward Richer Models of Rational Action

An alternative reaction to problems with the complexity of decision-theory, is to pursue richer, more reflective decision-theoretic approaches. This has been our approach. As the term bounded rationality has primarily been used to refer to heuristic approaches, several of us have used the term bounded optimality to refer to the pursuit of rationality under scarce resources through theoretically sound means.

The normative theory of rational beliefs and actions defined by the axioms of probability and utility has been the gold standard in many disciplines for dictating the behavior of an ideal agent. However, the limited scope and applicability of this theory is underscored by recent research on the construction of autonomous agents immersed in complex, competitive environments. Uncertainty about the value of problem solving and alternative physical actions often plagues relatively simple agents immersed in complex situations.

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. We seek to optimize the expected utility of an agent's behavior given a set of assertions about a system's state of knowledge about the problems the agent may face, about the expected costs of reasoning, and about its capabilities. In the general case, we must consider such knowledge as uncertain knowledge.

Some of the earliest discussion on the explicit integration of the costs of inference within the framework of normative rationality was by Good [2], who made the distinction between what he referred to as *Type I* and *Type II* rationality. Good defined Type I rationality as inference which is consistent with the axioms of decision theory without regard to the cost of inference. Type II rationality refers to behavior that includes the costs of the effort of reasoning. Good proposes Type III rationality—reasoning about the costs of reasoning at the Type II level. Good generalized the notion to Type $N$ rationality where $N$ refers to the particular level of analysis. After defining these classes of rationality, Good did little to enumerate a number of problems associated with subscription to the richer notions of rationality under resource constraints.

Although it is straightforward to describe such a probabilistic problem-solver in the abstract, critical problems arise in attempting to automate rational behavior. We will focus, in the proposed paper and presentation, on issues that challenge classical notions of rationality in the context of scarce resources and uncertain challenges. These fundamental problems motivate the extension of the classical model of rationality.

## 2.3 Goals of Research on Limited Rationality

Often the pursuit of bounded-optimal behavior is countered by reaction from the more traditional halls of AI, that there is no need to worry about rationality at all, that approximations and heuristic methods perform adequately. However, it is clear that, in many applications, rational approaches, by definition will be equivalent to, if not better than the ill-characterized approaches. Even small changes in performance can be important. Relative minor changes in performance can be crucial in arenas with high-stakes decision making such as medicine, in situations where it is important to consider behavior over a long period of time, and in competitive situations, where, in the long term slight advantages can accrue.

In addition to seeking better performing reasoners, we wish to elucidate general principles of reasoning under scarce resources. Attempting to construct rational creatures has already given us intuitions about the mental world of effective bounded-resource agents.

## 2.4 Problem and Opportunity Areas

There are opportunities for the development of flexible reasoning and representation strategies. There are a set of problems with the rational control of inference that are currently wide-open. We'll get to those in a few minutes. Other problem and opportunity areas include the compilation of reasoning verses dynamic computation of results, issues surrounding distributed and parallel problem solving, the dynamic reformulation of models, and problems with characterizing preference, performance and environment.

# 3 Components of Problem-Solving Value

We view the utility of a reasoner to be a function of the status of inference-level and object-level attributes that may or may not be represented. Object-level attributes are dimensions of a result or outcome associated with the value of action in the world. These attributes include such dimensions as accuracy and precision of a result. We use u sub o to represent object-level utility. We say that object-level value is a function of a multiattribute vector v that describes a result.

We use $u_i(\vec{r})$ to refer to the inference-related cost. This component is a function of a vector $\vec{r}$ representing relevant attributes of computational cost. Such attributes include dimensions of computational resource such as

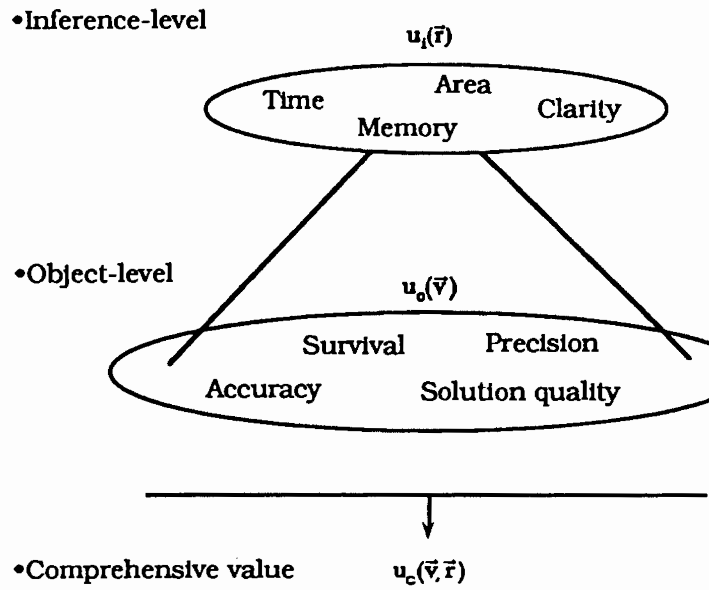


Figure 1: The utility of an agent in a context of challenges can be decomposed into object-level, inference-level, and comprehensive utilities.

the area of a VLSI chip, the available memory, or most commonly the time, it takes to generate a result. The object-level value, $u_o(\vec{v})$, is a function of a multiattribute resource vector $\vec{v}$. These object-level attributes include the cost

We say that the comprehensive utility, $u_c$, associated with an agent in the context of a problem is a function of the object-level and inference-level attributes. There is almost always uncertainty in the object level vector 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.

$$u_c = \sum_{\vec{v}} u_o(\vec{v}, \vec{r})\, p(\vec{v}|\vec{r})$$

It is often difficult to reason about the comprehensive utility directly. It can be useful to reason about the incremental changes in a result or action given additional expenditures of resource. Thus, the value of continuing to compute–or the value of computation is the change in the comprehensive utility. This is just the increase in object-level utility minus the cost of the additional computation.

$$\Delta u_c = \sum_{\vec{v}'} u_o(\vec{v}', \vec{r}')\, p(\vec{v}'|\vec{r}') \;-\; u_o(\vec{v}, \vec{r}) - [u_i(\vec{r}') - u_i(\vec{r})]$$

# 4 Flexible Inference and Representation

Let us first explore opportunities for innovation with the generation of flexible strategies. There are often uncertainties and variations in the amount of resource, such as time available for computation – thus, in general, strategies that are relatively insensitive to small variations in resource can be extremely valuable. A couple of years back, we enumerated several desired properties of strategies for reasoning under scarce resources. These include properties of flexibility. Two such properties are value continuity and value convergence. Value continuity refers to the notion that we wish strategies to generate results with continuous degrees of object-level utility with the expenditure of resource. Tom Dean and Mark Boddy have also highlighted the strategies exhibiting these properties. Value convergence says that we desire, continuous response strategies to converge on a result with optimal object-level value, given the availability of a a quantity of resources that is required for our best available method for completely solving a problem.

For example, classical normative reasoning has been based on a single model constructed as a static basis and acted upon by a single inference strategy. Simple examples demonstrate the value of reformulating a decision problem, considered ideal in a world of abundant resources, for situations of limited resources. In the general, the difficulty of a decision problem, the value of solving the problem, and the costs of computation may vary greatly. We have found tremendous opportunity in the area of reformulating traditional all-or-nothing approaches to flexible strategies. Several of you are familiar with our work on probabilistic inference and on sorting. We will review several examples from recent research for additional motivation.

## 4.1 Desiderata of Bounded-Resource Computation

Previous work has focused on elucidating desirable properties of problem-solving under ranging limitations in reasoning resources[3][1]. As an example, we wish to implement representation and inference methodologies that allow the most relevant processing to occur early on. Also, as many real-world applications involve reasoning under large variations in the amount of time available for inference, it is desirable to design inference strategies

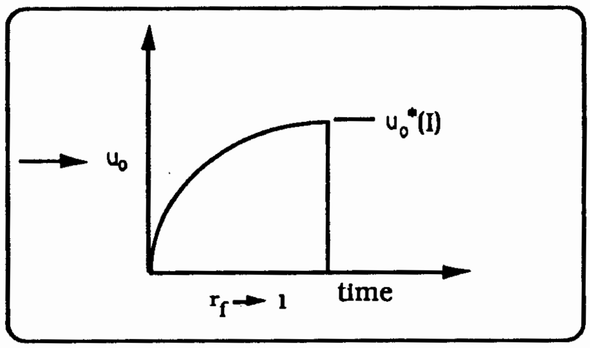


Figure 2: Strategies that show a continuous response in object-level value are useful for reasoning in contexts of uncerain and varying resource limitations.

that are insensitive to small changes in the amount of committed computation time. In particular, the development of flexible strategies that generate customized approaches over a wide range of resource availabilities is crucial. Such flexibility can be especially useful in light of uncertain deadlines and challenges. We will discuss fundamental aspects of flexiblity and will describe recent work on representation and inference strategies that configure knowledge and processing in a manner that allows for flexible responses to a range of computational resource availabilities[5].

1. 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 it ranges from zero to one. That is,

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

where $r_f' > r_f$ and $V_c(S, P, r_f, \xi)$ is the expected value of computation associated with an agent's applying inference strategy $S$ to problem $P$ with resource fraction $r_f$. 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*.

2. Convergence. We desire strategies that demonstrate convergence on the optimal object-related value as the resource fraction approaches one.

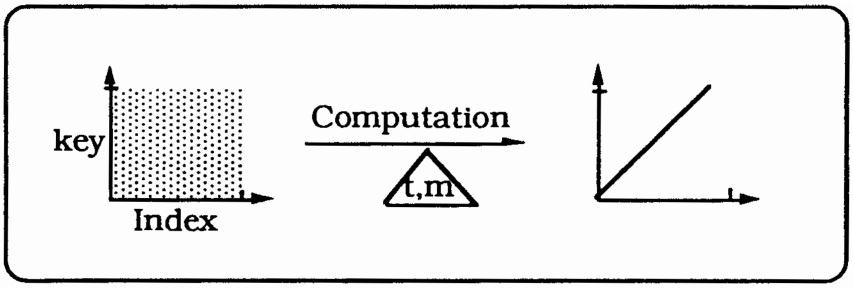


Figure 3: A representation of the problem of sorting a file of records that highlights the traditional all-or-nothing approach to computational problem solving.

That is, the object-related utility of a reasoning method should gracefully revert to ideal object-level rationality as the resources converge on $r_c$.

$$\lim_{r_f \to 1} V_o(S, P, r_f, \xi) = [V_o](P)$$

### 4.2 Multiple Dimensions of Partial Results

The Protos/algo project has explored the value of flexible reasoning with some classic computer science problems. This figure represents the traditional approach to sorting. On the left-hand side, we represent a randomly mixed file of records as a cloud of points in a two-dimensional space. The x-axis represents the file position and the y axis represents the sorting key. Sorting algorithms, fueled by time and memory, convert such a disheveled array into a final sort, with each item in its correct position. However, under varying and uncertain resource constraints we may not have time to complete our computation.

Several existing sorting algorithms generate different patterns of partial results with the expenditure of resource. For example, on the left we see the typical pattern of selection sort; on the right, the pattern of refinement of Shellsort is portrayed.

As we reported at AAAI last year, we can tease out useful dimensions of object level value in a partial sort, and can build multiattribute utility functions that tell us how much different partial sorts are worth to us. We can use this information in determining which algorithm to apply and how long to apply it before halting in different contexts.

An interesting highlight of this work was the demonstration of different

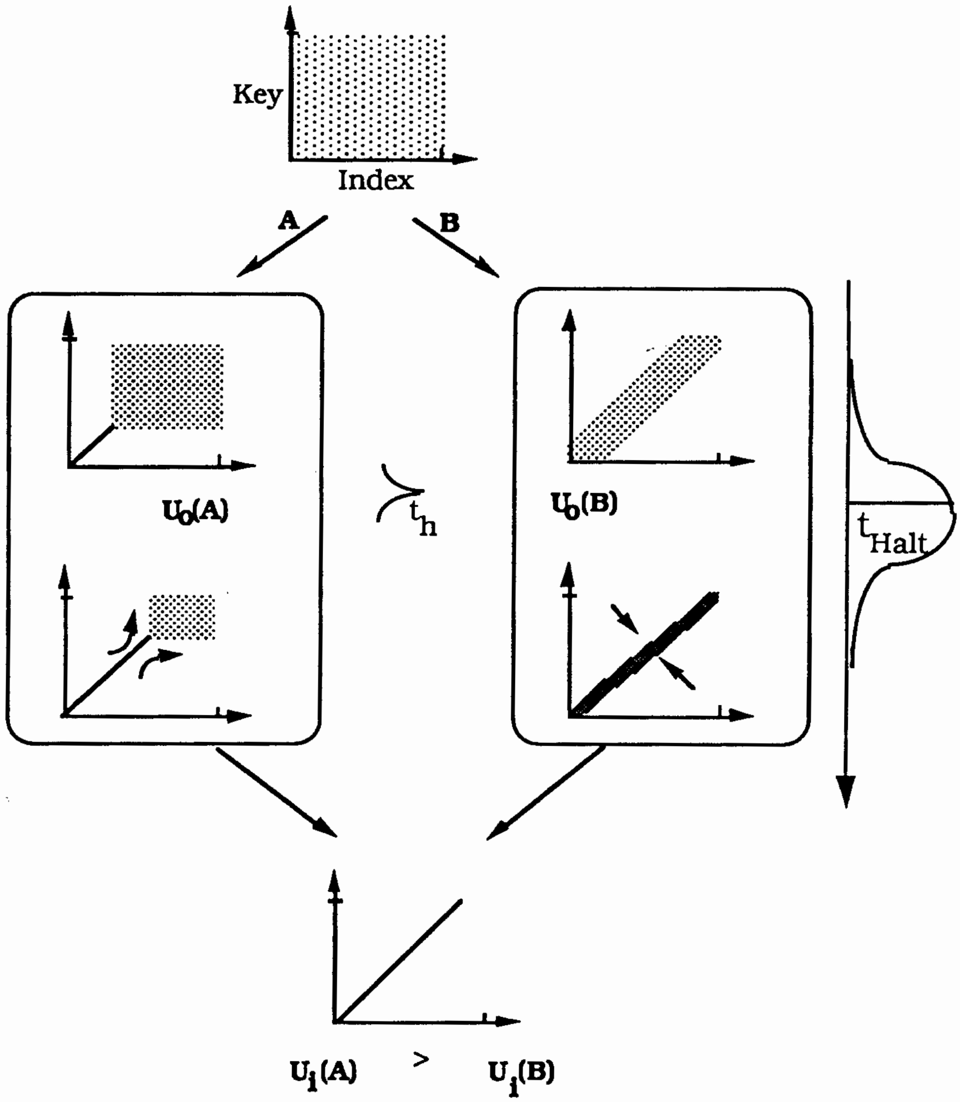


Figure 4: Different sorting algorithms produce different patterns of problem-solving behavior over time. The selection sort is pictured on the left. Shellsort is shown at right.

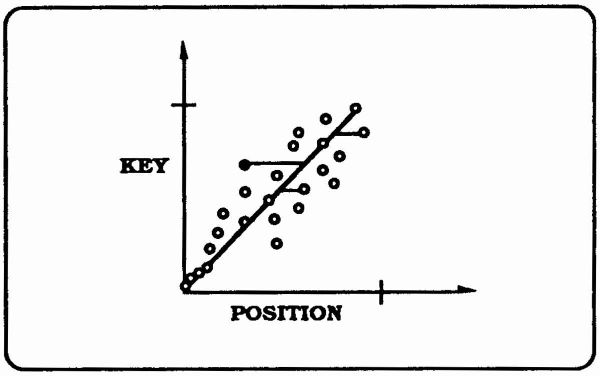


Figure 5: By focusing on human preferences, we can reason about the relative value of different partial results. This figure shows us several components of a partial sort.

problem solving trajectories that can be taken through a multiattribute space. In this picture the xy plane represents two attributes of a sort. The z axis represents time. Different algorithms fly different paths through this space.

The elucidation and characterization of different problem solving trajectories can be extremely useful in real-world problem solving, such as robot planning.

## 4.3 Object-level Innovation through Reformulation

### 4.3.1 Modulating completeness of problem solving

Although the sorting application has been instructive, we are most interested in decision theory under resource constraints. There is much opportunity for the development of well-characterized flexible strategies. Here is an influence diagram representation of a decision problem that focuses on the costs and benefits of placing a patient that is showing some signs of breathing difficulty on an automated respirator. In an influence diagram, and in decisionless belief networks, nodes represent propositions of interest and arcs represent dependencies among belief in the nodes. In our medical decision systems, we often must deal with large, multiply connected networks. We know that the complexity of inference within belief networks is NP-hard.

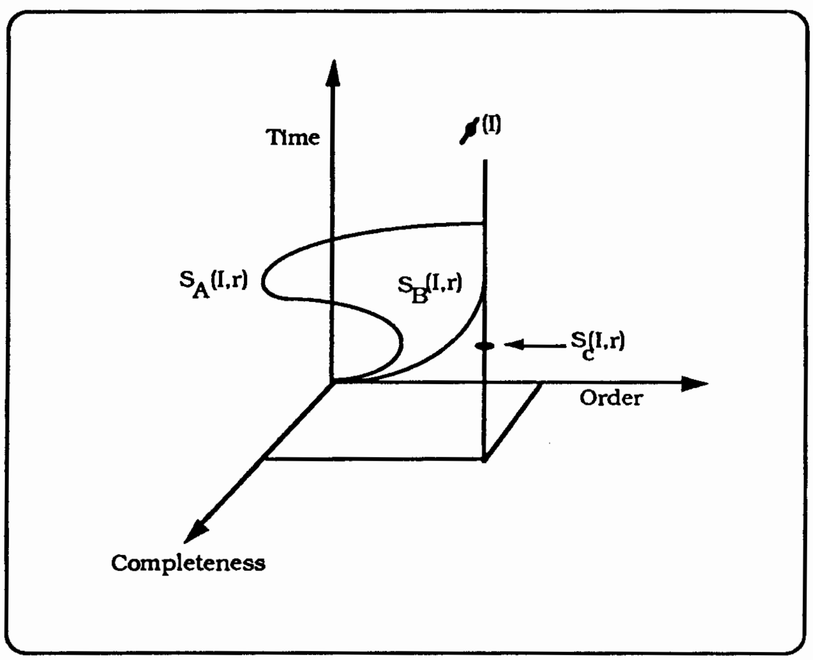


Figure 6: Different algorithms take different trajectories through a multiattribute space representing different dimensions of value of partial results.

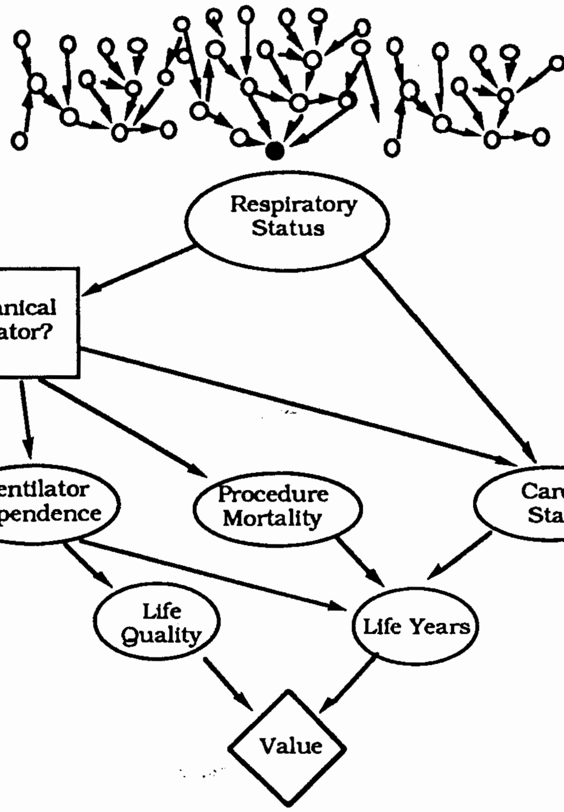


Figure 7: A representation of the decision of whether to place a patient on a respirator. Often, we must calculate the belief in a relevant proposition, such as the respiratory status of patient, by applying a complex belief network.

**Example: Bounded cutset conditioning** Here is a belief network, representing distinctions and relationships of importance for reasoning about a patient's health status in an intensive care unit. This was constructed by Ingo Beinlich, a doctoral student in our program, who also happens to be an experienced intensive-care-unit physician.

Just as we can apply sorting algorithms that generate partial sorts, we can design and apply strategies for doing flexible inference. As opposed to generating exact probabilities about a proposition of interest, we seek to generate distributions or logical bounds on a probability. We have worked on the generation of new types of flexible inference strategies, including those that are centered on the modulation of completeness and abstraction.

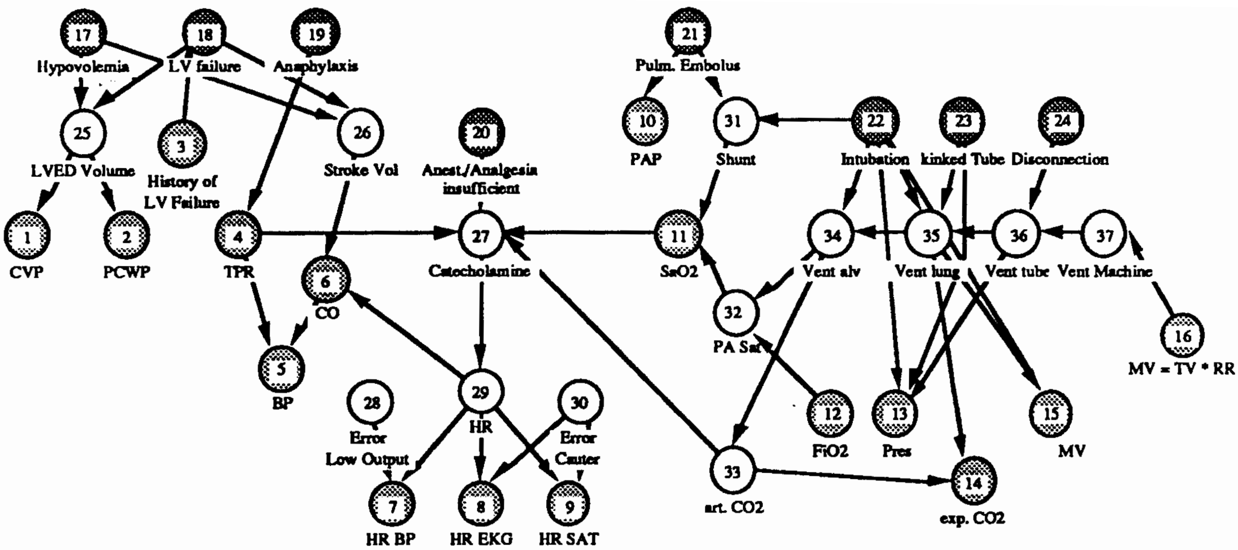


Figure 8: A belief network representing relevant distinctions, and probabilistic relationships in the intensive care unit.

As an example, we have examined the gracefulization of cutset-conditioning through modulating the completeness of analysis. In Pearl's method of cutset conditioning loops, in a belief network, are broken by a cutset, creating a set of subproblems or instances. All instances are solved to generate a final probability. The number of instances that must be solved are 2 to the size of the cutset. This can be quite large.

In bounded-cutset conditioning, we sort the instances by importance, and solve the problem sequentially. A bounding calculation reports the final bounds on a probability of interest as instances are solved. Here is an example of how the algorithm performs on the ICU network. We found that the bounds decay at rate that can be described with a negative exponential. This allows us to reason about the decay and half-life of bounds. This is useful in our work on the control of inference that we'll describe below.

### 4.3.2 Modulating abstraction of problem solving

Another approach to generating flexible inference is through modulating the level of abstraction at which reasoning occurs. In practice, a set of distinctions, viewed as optimal, is reformulated into more abstract entities at progressively greater levels of abstraction.

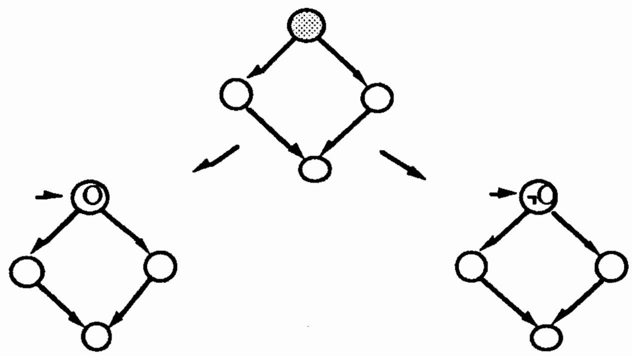

Figure 9: The application of the method of conditioning is based on the breaking of dependency cycles through determining a cutset that converts the belief network inference problem into a set of singly connected network subproblems.

David Heckerman and I have applied this abstraction modulation in the Pathfinder expert system for tissue pathology diagnosis.

Pathfinder uses such hierarchies of abstraction to control the level of detail at which value of information decisions are made.

There is much opportunity for the development of flexible inference strategies. This work focuses on the identification of continous dimensions of solution value, the the generaion of strategies that can refine results along these dimensions. The pursuit of optimal flexible strategies will also undoùbtely capture fundamental tradeoffs between object-level value and applied resources.

# 5 Rational Control of Problem-Solving

Let us now move to explore problems and opportunities with metareasoning. Given a set of approximation strategies, and varying costs of resources, we wish to rationally choose a strategy or set of strategies, and reason about how long to apply it. Some of the fundamental problems of metareasoning are the same, whether we are trying to control sorting, searching, or decision-theoretic inference. These problems pose exciting areas of research.

Applying a portion of available reasoning resources to consider the utility of alternative inference strategies or the value of continuing to refine a result before acting enables a computational agent to generate customized

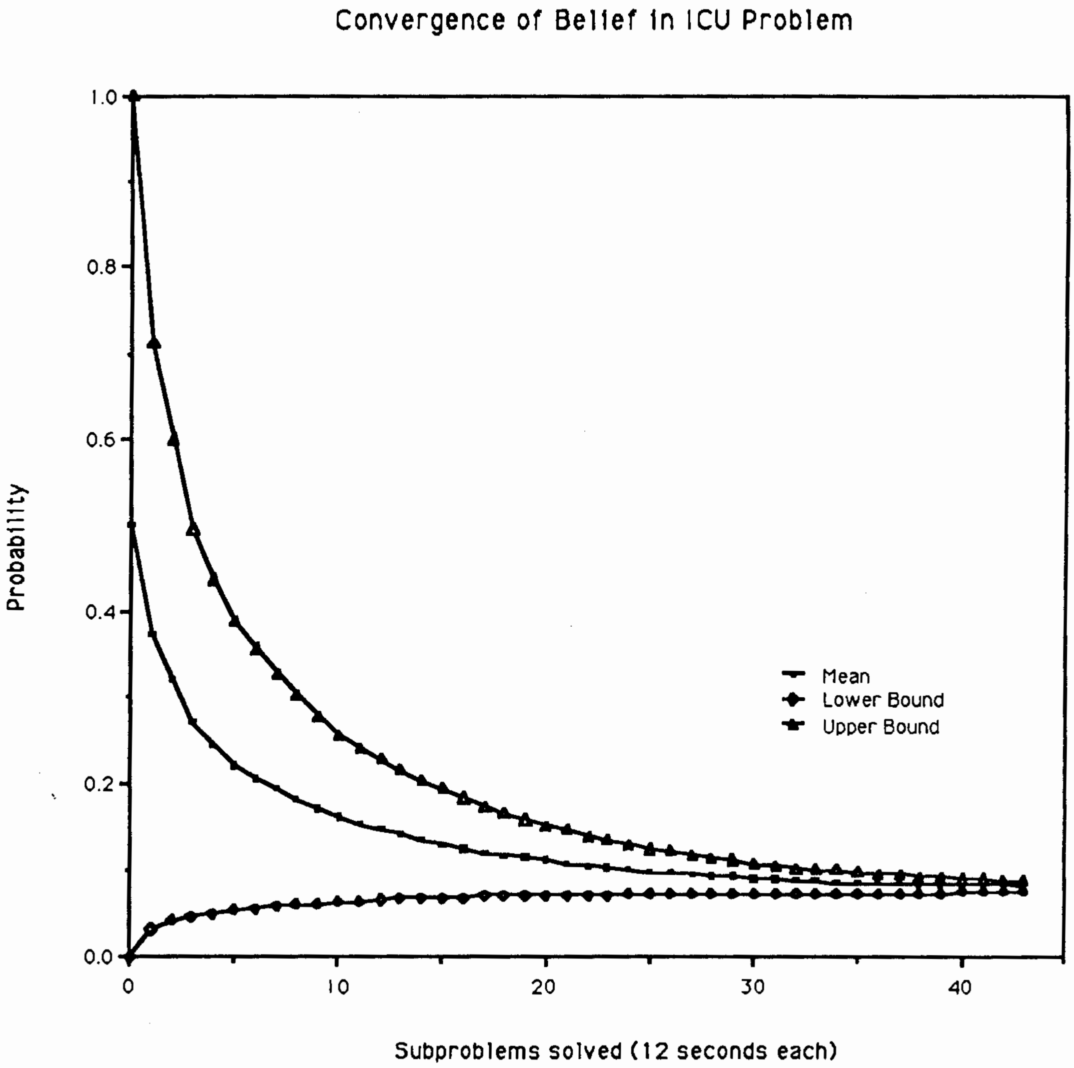


Figure 10: The peformance of bounded cutset conditioning on a specific updating of belief in the intensive care unit network. Convergence of the upper and lower bounds is pictured. The middle curve represents the mean at each point in the updating process.

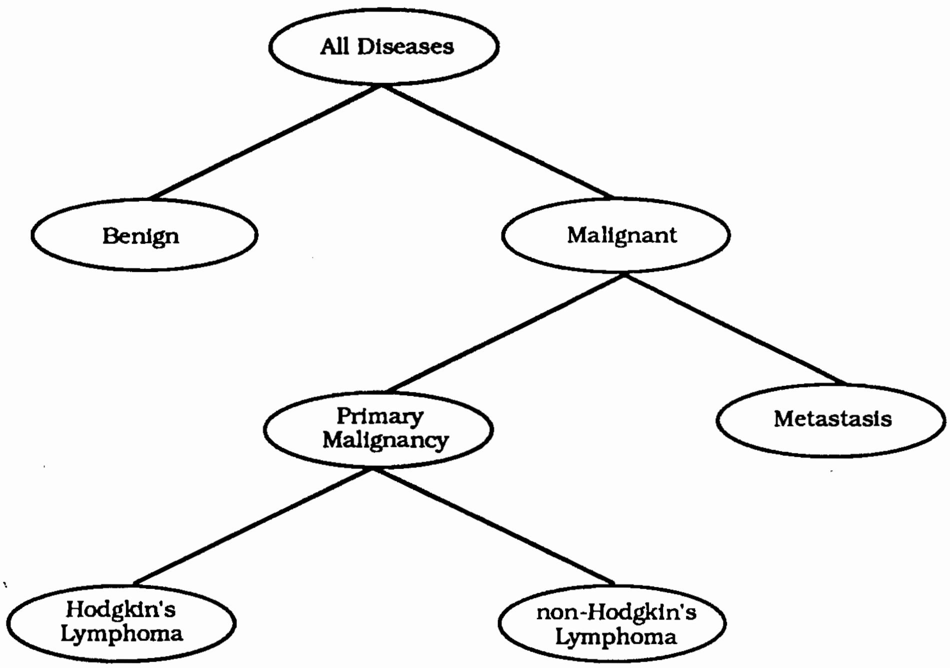


Figure 11: An abstraction hierarchy employed by physicians in reasoning about hematopathology diseases. The hierarchy represents a descent from the most abstract to more detailed classes of disease hypotheses.

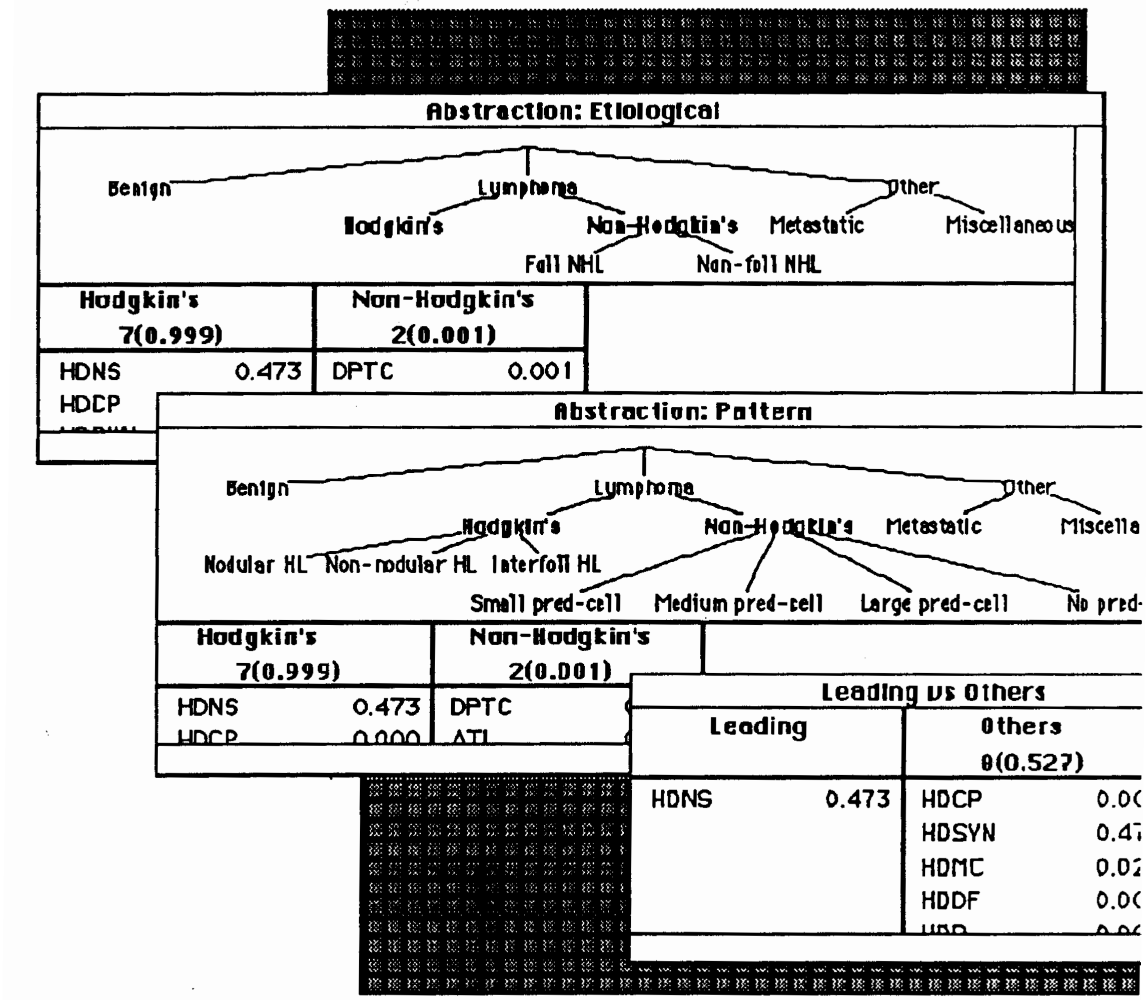


Figure 12: The Pathfinder system allows a user to probe a problem from different perspectives by making available different abstraction hierarchies.

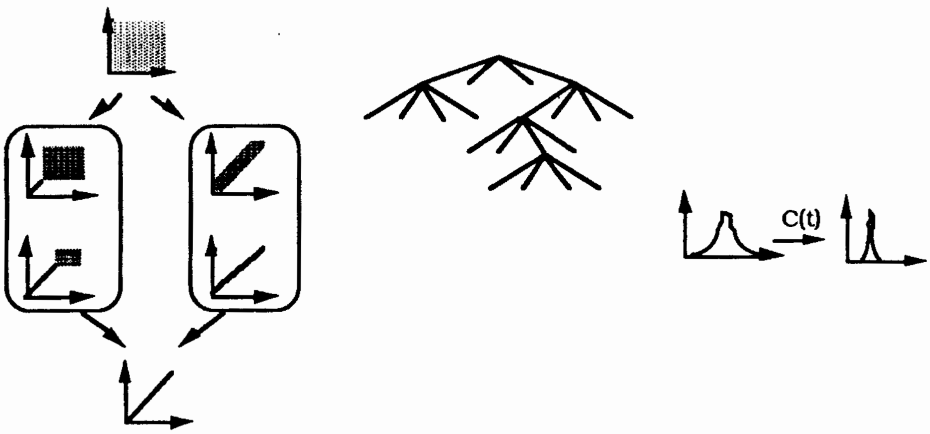


Figure 13: Fundamental principles and problems of controlling problem solving are the same, whether we wish to sort a file of records, search through a tree of alternative actions, or wish to perform probabilistic inference.

approaches to a wide variety of problems and time pressures. Such flexibility can be especially useful in light of uncertain costs and challenges. There has been recent work on the application of decision-theoretic metareasoning to decision-theoretic inference itself[3], as well as to game-playing search[8] and sorting[4].

## 5.1 A Model of Normative Reasoning Under Scarce Resources

Let me first introduce our work on the control of decision-theoretic inference. Our model of normative rationality is the application of decision theory to control decision theory itself. Thus we have attempted to build tractable representations of the control problem. We wish to reason about the costs and benefits of applying alternative strategies given uncertain knowledge about inference in an object-level decision problem.

We have worked to build an architecture that has a strategy base, containing knowledge about the expected performance of different reasoning strategies to refine different attributes, given different problem instances, a utility base that contains inference information about the combination of attributes into object-level utility, and a resource base that contains knowledge about the resource costs. Let's focus on the goal of the strategic reasoner.

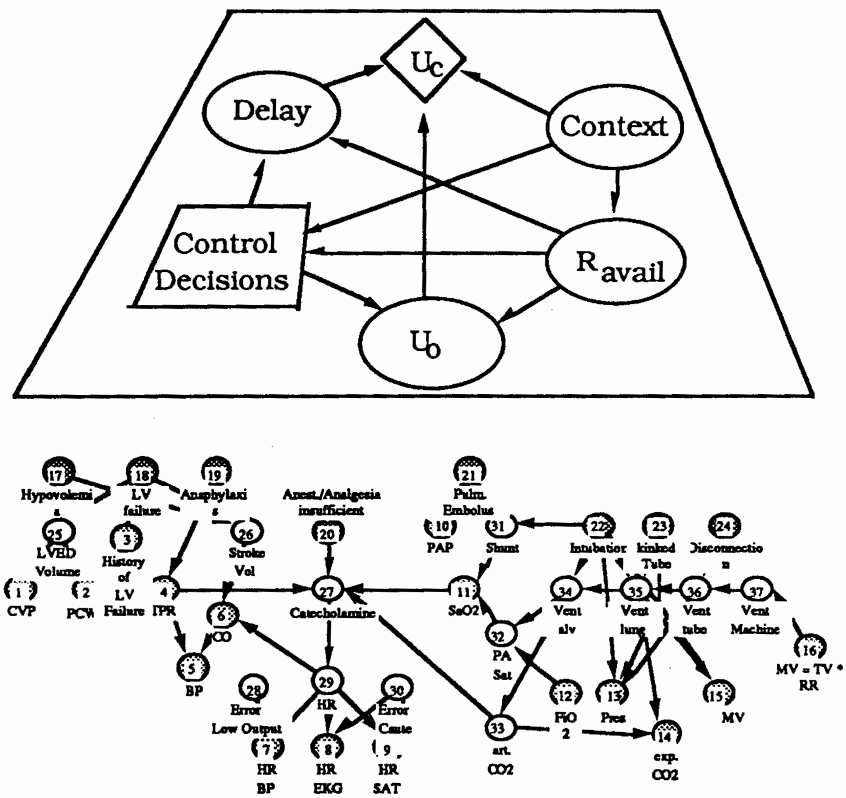


Figure 14: A model of normative inference, composed of an object-level inference problem and a control problem that focuses on the costs and benefits of alternative types and extents of object-level inference.

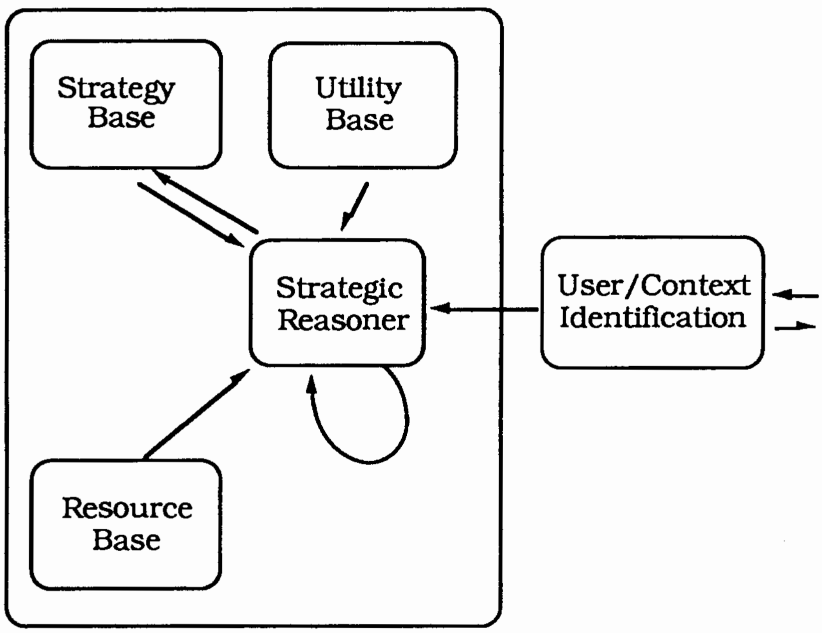


Figure 15: Fundamental components of an architecture for bounded-resource problem solving, centered on the separation of object-level performance, inference-related costs, object-level preferences, and a strategic reasoner that makes use of these classes of knowledge.

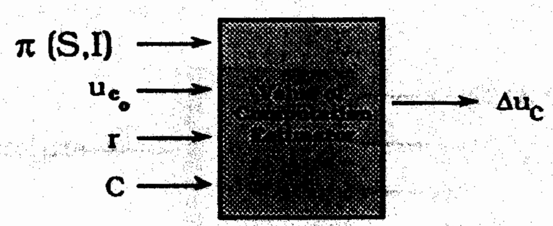


Figure 16: A central task in the implementation of decision-theoretic control is the construction of a mechanism for reasoning about the value of additional computation.

## 5.2 Valuating Future Computation

Our reasoner must determine the value of applying a strategy. Our approach, and almost all other approaches to date, have almost invariably sought to simplify the metalevel analysis by reasoning about the incremental benefit of applying a strategy. The problem is simplified by reasoning about the value of computing for another x seconds, and the identification of models that enable us to generate closed-form solutions to estimation of the value of computation.

Specifically, the task is focused on the construction of value of computation estimators that take as arguments, the current context, the amount of proposed reasoning resource, the current comprehensive utility, and a parameter that partially characterizes problem-solving performance. Such parameters serve as the "evidence" on the value of additional computation.

In the case of decision-theoretic inference, we tell our value of computation estimator about the decision problem, the object-level utility function, the current comprehensive value, and a parameter that represents belief about future belief, given a strategy and belief network instance.

This approach, and similar approaches, for the sake of tractability, makes a myopic, or greedy assumption. Our system only looks ahead to the next phase of computation.

## 5.3 Toward Richer Characterizations of Future Computation

There is also opportunity for elucidating, collecting, and applying richer predictors of the value of computation. In most work, we have used relatively simple measures such as the convergence of a probability or a distribution over utility.

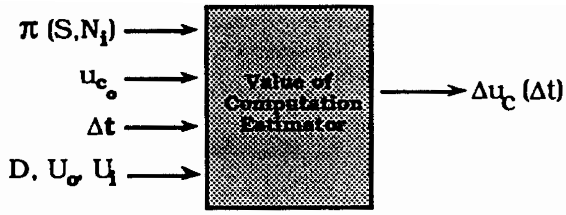


Figure 17: For the valuation of additional computation in reasoning about problems of belief and action, we must consider the current problem, the object-level, and inference-level utilities, and a parameter capturing the future belief that will be calculated with additional computation.

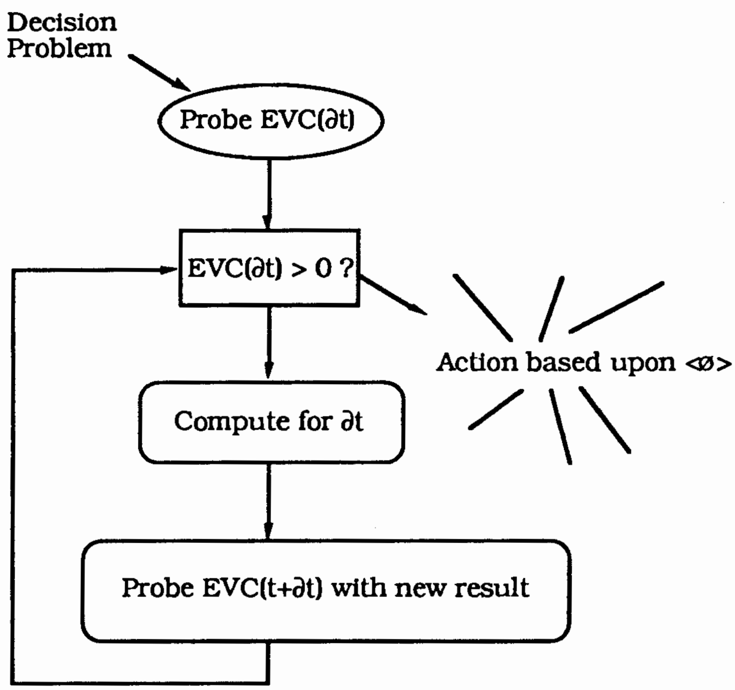


Figure 18: Most strategies explored to date make a myopic assumption in reasoning about the value of future computation.

In many cases, it will undoubtedly prove useful to collect problem-solving trajectories and representing long-term trends. However, in cases, where we have an explosion of the number of trajectories, we can abstract the information into classes–and build value of computation estimators that are hungry for this type of knowledge. Such characterizations of the performance of classes of reasoning strategies on classes of problems may be more useful than using the simpler parameters.

There are other classes of evidence about performance. There is opportunity for collecting knowledge about the value of alternative configuations of problem solving–for example, the value of adding another metalevel in a situation. Another kind of evidence is conditioned on current problem-solving progress. Useful intuition about how future computation will go might be gleaned from the recent history of progress.

### 5.3.1 Toward a Tractable Economics of Computation

There is quite a bit of research opportunity at moving to more global analyses. This work involves encoding expected solution trajectories. By making assumptions (hopefully valid) about the object-level utility with the expenditure of resource, and about the nature of the costs, we can do tractable global analyses. For example, if we have an object level utility function that is monotonically increasing and convex and a linear cost function, we can easily determine an optimal halting time.

In fact, we can apply well known theoroms from economics to reason about the production of utility for an agent, and optimal values for halting given changes in the criticality of a context.

We can also prove dominance of whole sets of strategies over other strategies with ease.

## 5.4 Controlling control: Reasoning About Analytic Regress

So far, we have assumed that we will be able to do tractable metareasoning. However, our attempts to simplify the metareasoning problem, as discussed above, will not always pay off. Let me turn to that age old problem of infinite analytic regress. If indeed, adding a metalevel is so useful, why not add another to assist with the first metareasoning problem? Where do we stop? We see this as a rich problem area for bounded resource inference-and an important one for generating truly bounded-optimal reasoners.

There is an urgent need for understanding the costs and benefits of increasing numbers of metalevels. Promising areas include bounding the num-

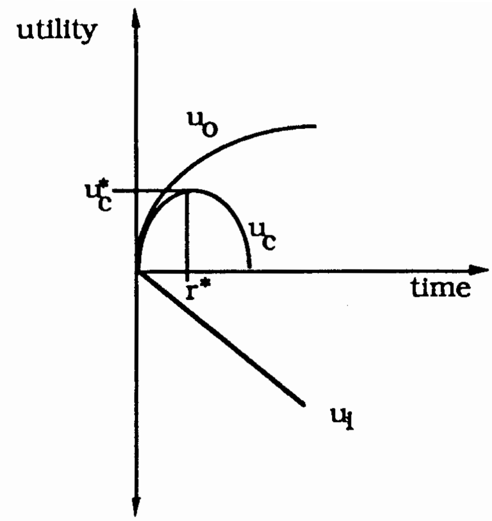


Figure 19: Notions of marginal production from work in economics allows us to reason more globally about the optimal quantity of resource to expend in solving a problem.

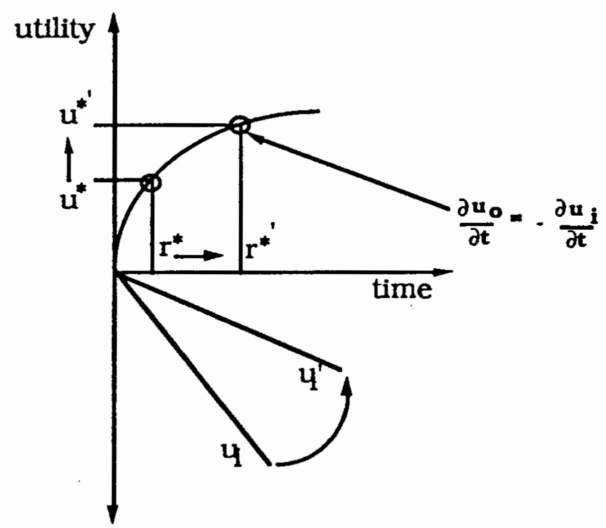


Figure 20: We can apply economic analyses to determine how changes in the criticality of a context will change the optimal quantity of resource to apply to a problem and the expected comprehensive value of computation.

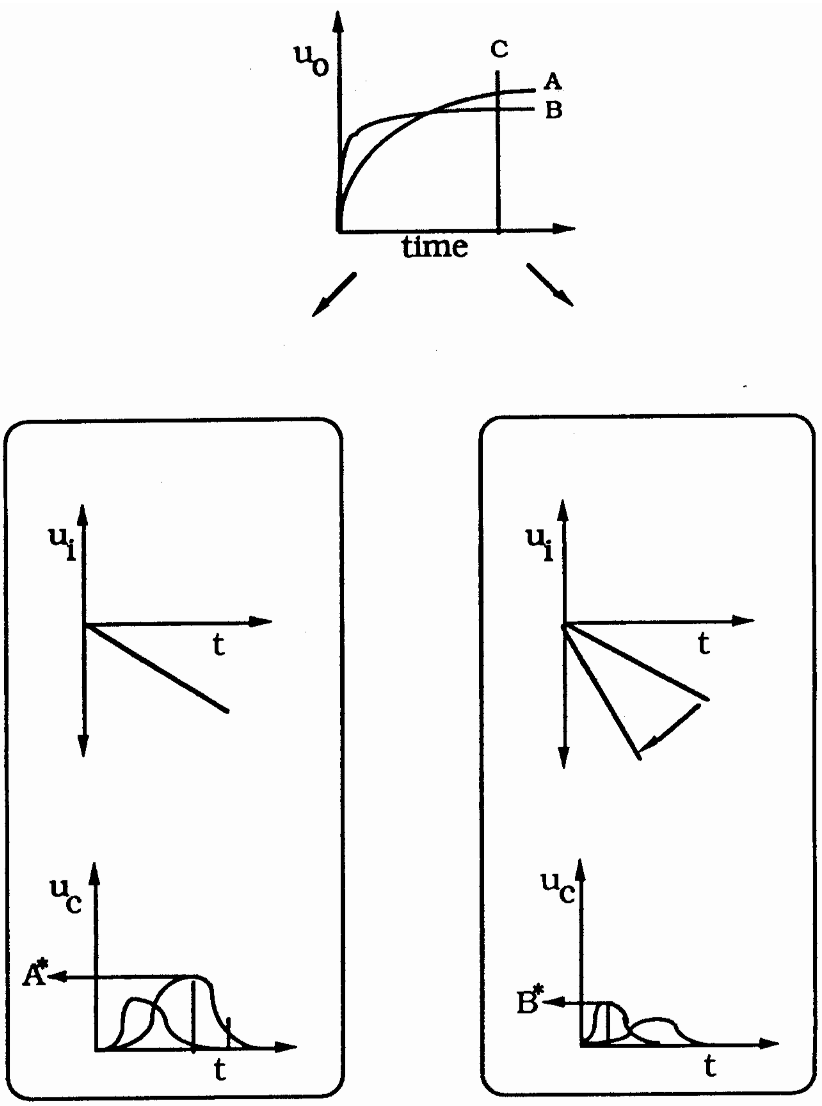


Figure 21: There is opportunity for the application of proofs of dominance in the selection of best strategy for the context at hand.

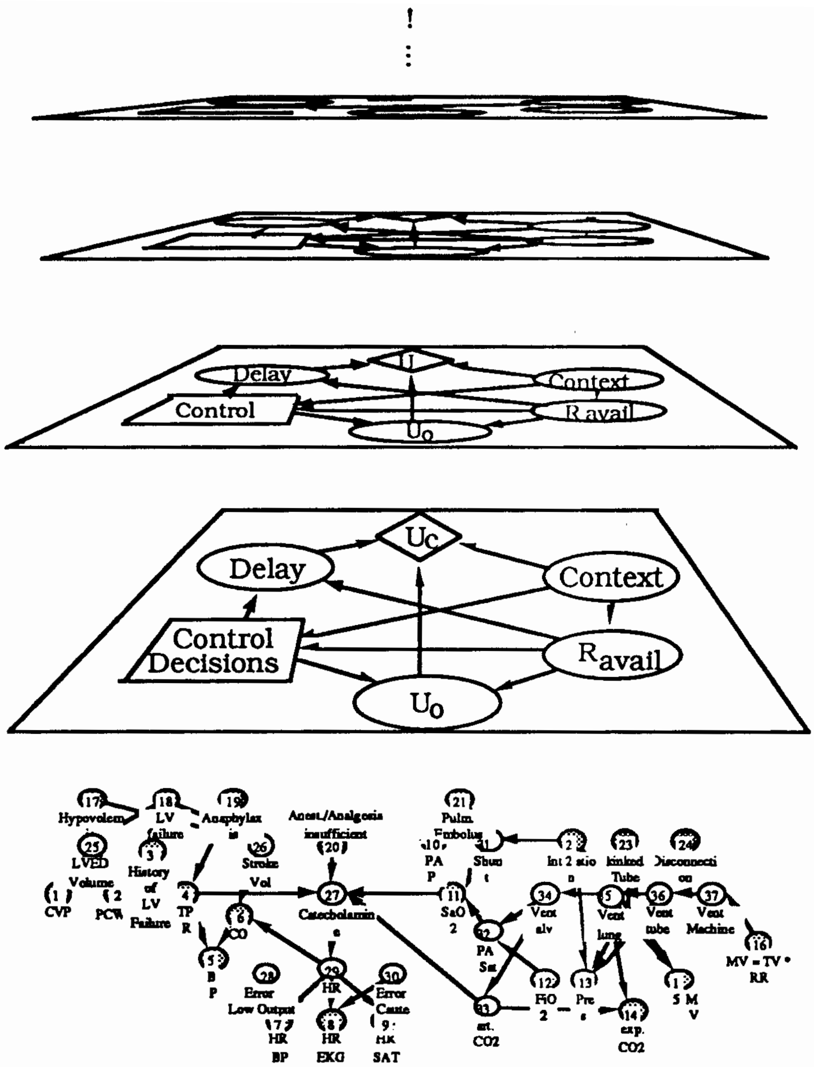


Figure 22: The problem of analytic regress becomes important in situations where the metalevel analysis of problem solving requires significant quantities of computational resource.

ber of metalevels and reporting a level of error on incomplete meta-analysis. Other areas of interest include performing sensitivity analyses on the worth of additional metalevels. As an example, in certain problem areas, and for specific metalevel configurations, we will be able to prove convergence on an estimate of the value of computation.

In summary, there is rich research ahead for moving beyond myopic analyses, for enriching the evidence used in value of computation reasoning, and in seeking sound approaches to analytic regress.

# 6 Reasoning about Decision Models

The process of constructing and valuating alternative decision models is a crucial part of decision making under limited and uncertain resources. Decision theory says nothing about the construction, selection, and updating of alternative decision models. There are different approaches to the construction and metalevel valuation of decision bases. We will describe several approaches. We will examine the difficulty of valuating alternative models arising in complex interactions between decision models and the preferred methods for solving specific models.

# 7 Dynamic Computation verses Compilation

We have only begun to study the rich interrelationships between the dynamic computation of results on the one hand and the use of default reasoning and compiled knowledge, on the other. The development of techniques and of general architectures for utilizing a spectrum of default and pre-computed results will have great payoff.

We often have quite a bit more time to solve a problem in the design phase, than we do when we immerse our agents in some chaotic enviroment. Also, in many areas, it is safe to say that an agent will spend most of its time in an offline state. Why not build our agent's to be anxiety-ridden critters that worry about expected problems – much as we mortals do.

Recent work by Herskovits has focused on the analysis of the usefulness of speeding up inference by building a tree of cached probabilities ahead of time. A tree of cases is constructed by doing simulation weighted by utility. We are now working to grant access of such knowledge to a metareasoner. The metareasoner will have access to a large tree that can be easily searched, and to knowledge about what can be found in the tree. If a search

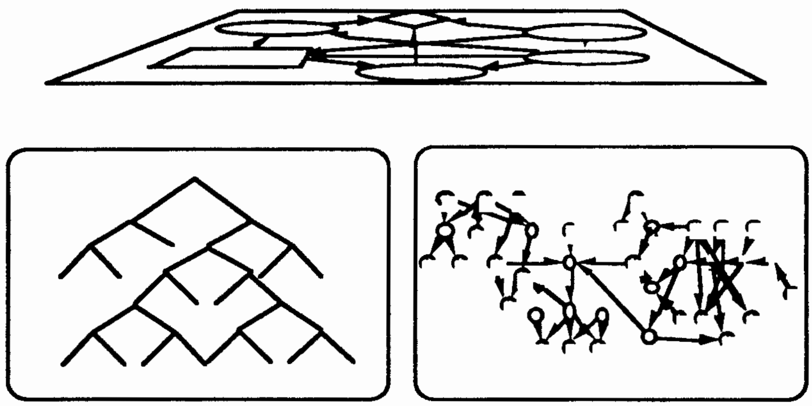

Figure 23: There is much opportunity for storing compiled belief or action and reasoning about the relative value of precomputation and caching verses computing a response to a problem challenge.

is unsuccessful, the reasoner can lick its wounds and begin to do inference. Given limitations in memory and time, a number of interesting questions arise about when to throw out a probability–or to restructure a tree. Thus, there are interesting research problems on the control of learning.

For example, if we store the probability that a probability or compiled action recommendation is stored in the cache, we can reason about the expected value of searching the tree for the probability. We just multiply the probability of finding the result by the utility of having the result at time (searchtime) and subtract the utility of computing the result with the expected time that it will take to do inference.

The savings in utility storing a compiled result in the cache verses performing inference can be used to control the original instantiation of the cache. We wish to first store rules or probabilities that we know will have utility greater than the cost of memory.

## 8 Concurrent and Distributed Strategies

There is also quite a bit of opportunity for innovation with parallel and distributed approaches. To date, we have worked primarily on a single processor model. That is, computation must be separately focused either on metareasoning, on obect-level problem solving, or on the performance of an action.

### 8.1 Concurrent Metareasoning, Reasoning, and Action

Concurrency for these tasks will undoubtedly have great payoff in many problem areas. Given the ability of a metareasoner to fundamentally change the entire course of reasoning, having it available at all times to monitor the problem-solving situation and to continue to reason, can have a great leveraging effect on performance. Also, there are opportunities for distributed control. Jon Doyle has been especially interested, –and has done interesting work in this area.

### 8.2 Object-level Innovation through Concurrency

Finally, there may be great payoffs for concurrent flexible inference. We have experimented with concurrent problem solving with the bounded cutset conditioning.

In this case, we show the effect of concurrent analyses with two different cutsets of a problem with the ICU problem. We could see that with two bounding algorithms running side by side, that we get faster tightening of the bounds.

The last area we'll address is that of problems with preferences and performance. Understanding the structure of utility is central to reasoning about the optimality and relative optimality of agents in different contexts. It is clear that as we move from simple "one-shot" decisions, as Mike Wellman has referred to those gems of common decision analysis, to long-term behavior, that it becomes more difficult to characterize performance. Questions about what it is that our agents, or that we, as their clients, are trying to optimize arise.

## 9 Preference, Performance, and Environment

The summarization of our preferences about complex histories of behavior ahead of time can be extremely difficult. There has been extremely little work on the valuation of anything except abstract "endpoints" of analysis.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. It is important to come up with models of utility summarization so that we design and reason about the relatve performance of our agents.

The determination of true bounded optimality requires proving tight

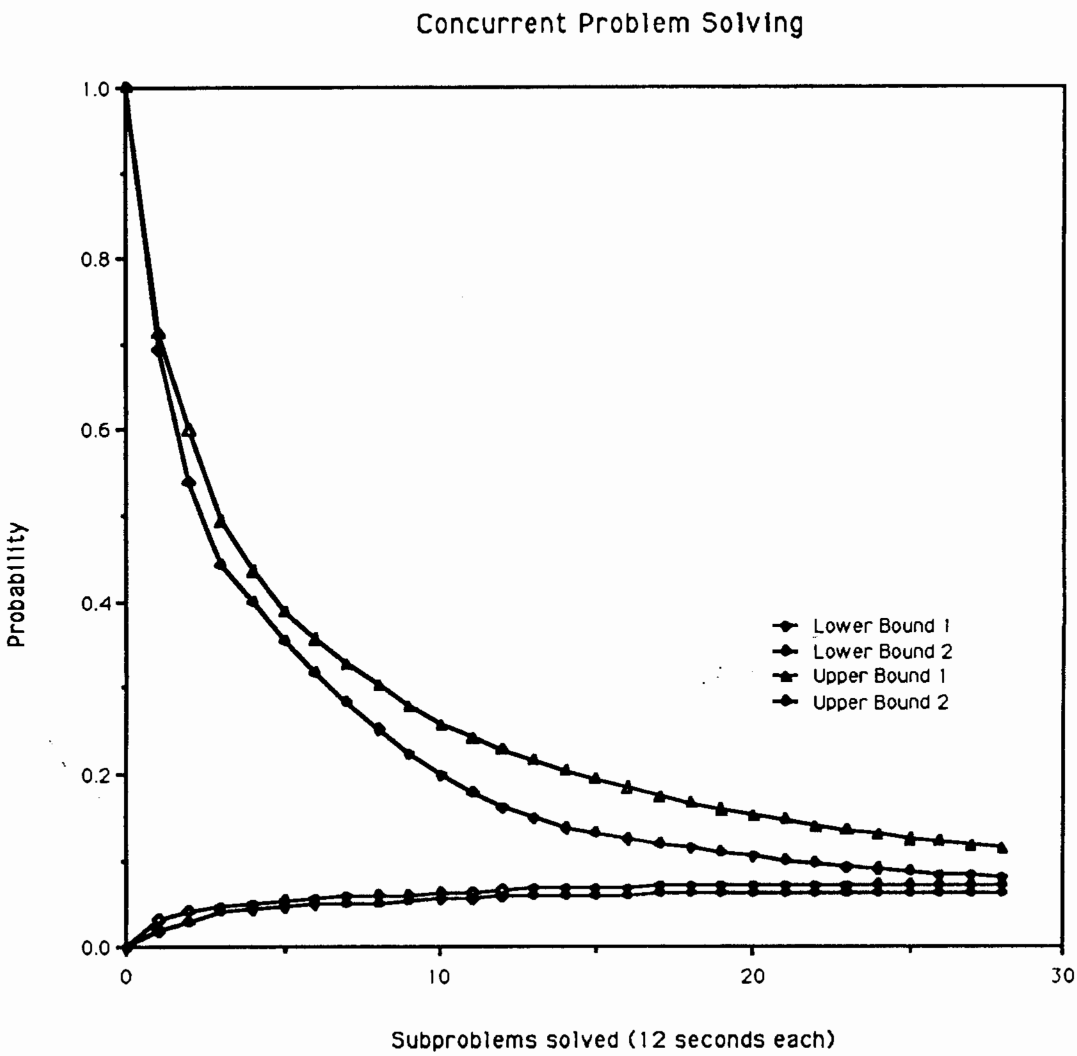


Figure 24: The concurrent application of two bounded cutset conditioning analyses.

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.

So to reason about the relative optimality of different problem solvers, we will need to characterize–and agree upon our characterizations of utility and environment. We have worked a bit on this, for example, proposing an independent challenge model where the relative optimality of a robot is gauged in the context of the distribution and frequencies of independent challenges over time.

As an example, we shall define a form of bounded optimality that is framed in terms of an agent limited to a distinct set of reasoning strategies. This statement of *bounded strategic optimality* is relevant to current research on the solution of probabilistic inference problems with alternative inference 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,

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

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

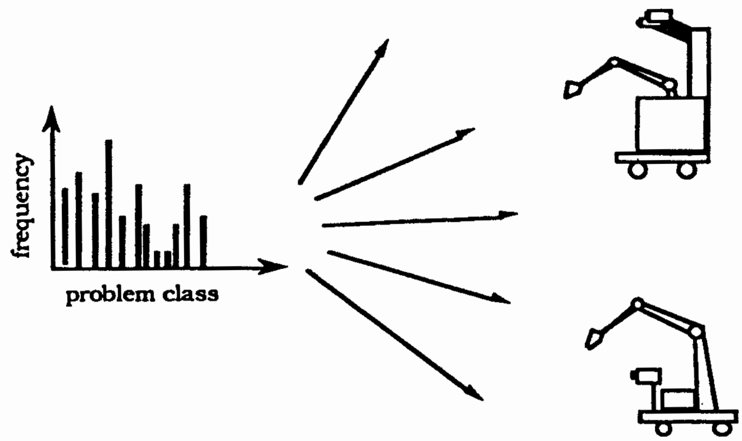


Figure 25: Reasoning about the expected value of different problem-solvers requires the characterization of the environment in which they are to be immersed.

perspective can be useful in comparing the effectiveness of agents, with different compositions and abilities, immersed in distinct problem contexts. The value associated with the use of a reasoning system over a period of time $t$, $V_s(t)$, is

$$V_s(t) = \sum_{i=1}^{n} t \; f_{P_i} * V_c[S^*(P_i, \xi)] \quad - \quad V_h$$

where $f_{P_i}$ is the frequency of problem challenge $P_i$, $S^*(P_i, \xi)$ is the best computation strategy available to an agent, and $V_h$ is the cost associated with the hardware that comprises the reasoning system.[1] We have separated the expense of the commodity that is required to construct the problem-solver's hardware from computation-based resources such as time. In this *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. This description, and other definitions of agency preference, can be useful in reasoning about such factors as the relative performance of different agents in specific contexts,

[1] We are assuming no maintenance costs; inclusion of timewise maintenance costs can be represented as a cost function over time.

the value of learning in a domain, and the utility of adding a new capability to an agent's problem-solving repertory.

As agents of different constituencies may have different "best" default reactions before engaging an explicit computation strategy $S$, it can be useful to consider the available pattern of default responses that may be further refined with computation. Thus, we say

$$V_s(t) = \sum_{i=1}^{n} t\ f_{P_i} * [V_d(P_i, \xi) + \Delta V_c(S^*[P_i, \xi])] \ - \ V_h$$

where $V_d$ is the expected value associated with an agent taking default action in response to a problem challenge, and $\Delta V_c$ is the expected net value of the best strategy available to the agent, given the default response. Work on the compilation of beliefs and actions, as described above, can increase the value of the default responses, usually associated with additional expense of $V_h$ associated with larger memory caches.

**Value of Agency** It can be useful to consider the fundamental benefit of introducing an intelligent agent into some environment. We define the value of agency, $V_a(t)$, over a time period $t$, as the difference in the value of a system, $V_s(t)$, defined above, and the value of some simple non-agent policy, $S_\emptyset$, that would be undertaken in the absence of a reasoner, such as never acting, or acting randomly,

$$V_a(t) = V_s(t) - V_\emptyset(t)$$

where

$$V_\emptyset(t) = \sum_{i=1}^{n} t\ f_{P_i} * V_o[S_\emptyset(P_i)]$$

Although our goal is to prove bounded optimality, it can be important to compare relative optimality of two different agents. For example, given a set of agents, an environment $C$, and time horizon $t$, we may prefer to employ an agent, $A^*$,

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

where $f_{P_i}(C)$ is the frequency of problem type $P_i$ in context $C$; $V_d$ is the expected value associated with an agent taking default action in response to a problem challenge, and $\Delta V_c$ is the expected increase in value generated by the best strategy available to the agent.

nally, we mentioned that it is of crucial importance that we decide what it is that we are trying to optimize given multiple decisions in complex contexts. This will take reflection about our own utilities over time and through a history of challenges, as well as the construction of usable models that can mathematically represent such utility.

## References


[1] T.L. Dean and K. Kanazawn. Probabilistic temporal reasoning. Technical report, Brown University, May, 1988.

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

[3] E.J. Horvitz. Reasoning about beliefs and actions under computational resource constraints. In *Proceedings of Third Workshop on Uncertainty in Artificial Intelligence*, Seattle, Washington, July 1987. American Association for Artificial Intelligence. Also appearing in L. Kanal, T. Levitt, and J. Lemmer, ed., *Uncertainty in Artificial Intelligence 3,* North Holland.

[4] E.J. Horvitz. Reasoning under varying and uncertain resource constraints. In *Proceedings AAAI-88 Seventh National Conference on Artificial Intelligence*. American Association for Artificial Intelligence, August 1988. Also available as Technical Report KSL-88-35, Knowledge Systems Laboratory, Stanford University, April 1988.

[5] E.J. Horvitz, H.J. Suermondt, and G.F. Cooper. Bounded cutset conditioning: An incremental-refinement approach to belief under uncertain resources. Technical report, Stanford University, July 1988. KSL-88-36.

[6] R.L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs.* Wiley and Sons, New York, 1976.

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

[8] S.J. Russell and E. Wefald. Multi-level decision-theoretic search. In *Proceedings of the AAAI Spring Symposium on Game Playing*, Stanford, CA, March 1988. American Association for Artificial Intelligence.

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