Reflection and Action Under Scarce Resources: Theoretical Principles and Empirical Study

Eric J. Horvitz, Gregory F. Cooper, David E. Heckerman

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

Abstract

We define and exercise the expected value of computation as a fundamental component of reflection about alternative inference strategies. We present a portion of Protos research focused on the interlacing of reflection and action under scarce resources, and discuss how the techniques have been applied in a high-stakes medical domain. The work centers on endowing a computational agent with the ability to harness incomplete characterizations of problem-solving performance to control the amount of effort applied to a problem or subproblem, before taking action in the world or turning to another problem. We explore the use of the techniques in controlling decision-theoretic inference itself, and pose the approach as a model of rationality under scarce resources.

1 Reflection and Flexibility

Reflection about the course of problem solving and about the interleaving of problem solving and physical activity is a hallmark of intelligent behavior. 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 custom-tailored approaches to a wide variety of problems, under different time pressures. Such flexibility can be especially useful in light of uncertain deadlines and challenges. Uncertainty about problems and problem solving plagues simple agents immersed in complex environments. Constraints on an agent's reasoning and representation resources lead to inescapable uncertainties about the problems that may be faced and about the value of future reasoning in solving those problems.

The Protos project has pursued the use of decision theory for real-time control and offline problem-solving design. The work has highlighted opportunities for the principled control of reasoning under scarce resources with problems in sorting and searching and with decision-theoretic inference itself [Horvitz, 1987a]. We havp particularly dwelled on the decision-theoretic control of decision-theoretic inference as a model of rational computational inference under resource constraints. In this paper, we present a component of this work centering on the use of incomplete characterizations of the progression of probabilistic inference to reason about the value of continuing to reflect about a problem versus taking action in the world. This methodology uses knowledge that partially characterizes relevant dimensions of problem-solving performance. Such knowledge can be learned and refined with experience. We shall introduce components of utility for computational or real-world actions, and define the expected value of computation in terms of the likelihood of future probability distributions over the truth of relevant propositions about the state of the world. After discussing the theoretical principles and empirical results, we describe a component of research centering on the offline analysis of problem-solving trajectories. Such offline musing, weighted by expected challenges, can be important in real-time reflection about problem solving.

2 Decision-Theoretic Valuation

Decision theory provides the foundations for a principled approach to metalevel decision making under uncertainty. Decision-theoretic metareasoning can be especially useful in reasoning about the selection, and optimal halting time, of reasoning strategies that incrementally refine results as scarce resources are expended [Horvitz, 1987b, Dean and Boddy, 1988].

We usu *comprehensive value,* U_c, to refer to the utility associated with the value attributed to the state of an agent in the world. This value is a function of the problem at hand, of the agent's best default action, and of the stakes of a decision problem. We call the net change expected in the comprehensive value, in return for some allocation of computational resource, the expected value of computation (EVC). It is often useful to view the comprehensive utility, at any point in the reasoning process, as a function of two components of utility: the *object-level* utility, u_0, and the *inference-*

related cost, u_i.[1] The *object-level* utility of a strategy is the expected utility associated with a computer result or state of the world. We say that the object-level utility is a function of a vector of attributes, $\vec{v}$. For example we may assign an object-level utility to an incompletely sorted file of records, based on several different dimensions of incompleteness. The *inference-related* component is the sum of the expected disutility intrinsically associated with, or required by, the process of problem solving. This cost can include the disutility of delaying an action while waiting for a reasoner to infer a recommendation. In general, the inference-related cost is a function of a vector of resource attributes, $\vec{r}$, representing the quantity that has been expended of such commodities as time and memory.

There is generally uncertainty in the object-level state resulting from the expenditure of computational resources. Thus, in the general case, we must sum over a probability distribution of object-level attributes to generate an expected comprehensive utility. If $\vec{v}$ and $\vec{r}$ are the vectors representing object-level and inference-related attributes without additional computation, respectively, and the $\vec{v}\,'$ and $\vec{r}\,'$ are the revised vectors, expected with additional computation, the net, or change in, comprehensive utility, given some allocation of resources is

$$\mathrm{EVC}(\vec{r}\,') = \sum_{\vec{v}\,'} u_c(\vec{v}\,',\vec{r}\,')\,p(\vec{v}\,'|\vec{r}\,') \;-\; u_c(\vec{v},\vec{r})$$

In cases where the inference-related and object-level utilities can be decomposed, and are related through addition, the EVC is just the difference between the increase in object-level utility and the cost of the additional computation,

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

In another study, we considered the refinement of multidimensional attributes of partial results with computation [Horvitz, 1988]. Here, we will simplify our object-level focus to a probability of a state in the world, *H*, and the quality of an associated decision to act, *A*, given uncertainty about the truth of the state. We will simplify the inference-related component to a consideration of computation time.

The decision-theoretic approach to metareasoning in difficult machine intelligence problems was introduced by I.J. Good over 2 decades ago, in the context of the control of game-playing search [Good, 1968]. Good had earlier discussed the explicit integration of the costs of inference within the framework of normative rationality, defining Type 1 rationality as inference that is consistent with the axioms of decision theory, regardless of the cost of inference, and Type II rationality as behavior that takes into consideration the costs of reasoning [Good, 1952]. Related work in decision science has focused on the likely benefit of expending effort for decision analyses [Matheson, 1968, Watson and Brown, 1978]. Our group

More comprehensive notions of the value of a reasoning system in an environment are discussed in [Horvitz, 1987b].

researched the general applicability of decision-theoretic control of computation, with an emphasis on metareasoning problems with probabilistic inference and knowledge representation [Horvitz, 1987b]. Early investigation demonstrated that multiattribute decision-theoretic control of reasoning had promise for guiding the solution of a variety of tasks, including such fundamental problems as sorting a file of records or searching a large tree of possibilities [Horvitz, 1987a]. Indeed, there have been recent studies of the value of computation in the control of sorting [Horvitz, 1988] and of game-playing search [Russell and Wefald, 1988, Hansson and Mayer, 1989]. In related research on the control of logical inference, Smith, and Treitel and Genesereth, have explored the use of decision theory for selecting alternative logical reasoning strategies [Smith, 1986, Treitel and Genesereth, 1986].

3 Complexity of Inference

In reasoning about real-world actions under uncertainty, an agent generally must consider alternative decisions and outcomes, preferences about the possible outcomes, and the uncertain relationships among actions and outcomes. We have been investigating the use of *influence diagrams* [Howard and Matheson, 1981] for representing and solving automated reasoning problems. The influence diagram is an acyclic directed graph containing nodes representing propositions and arcs representing interactions between the nodes. Nodes represent a set of mutually exclusive and exhaustive states; arcs capture probabilistic relationships between the nodes. Influence diagrams without preference or decision information are termed *belief networks*. A belief network defines a model for doing probabilistic inference in response to changes in information.

The problem of probabilistic inference with belief networks is $\mathcal{NP}$-hard [Cooper, 1987]. Thus, we can expect algorithms for doing inference to have a worst-case time complexity that is exponential in the size of the problem (e.g., the number of hypotheses and pieces of evidence). Some methods for inference in belief networks attempt to dodge intractability by exploiting independence relations to avoid the explicit calculation of the joint-probability distribution. A variety of exact methods has been developed, each designed to operate on particular topologies of belief networks [Horvitz et al., 1988a]. Other methods forego exact calculation of probabilities; these approximation techniques produce partial results as distributions or bounds over probabilities of interest. The complexity of precise inference and the availability of alternative reasoning approaches highlight the need for robust approximation strategies and intelligent control techniques. We have sought to develop and control decision-theoretic inference for reasoning under uncertainty in high-stakes and time-pressured applications, such as medical decision making.

4 Decisions Under Scarce Resources

Let us explore concerns that arise in automated decision making under scarce resources. The graph in the lower
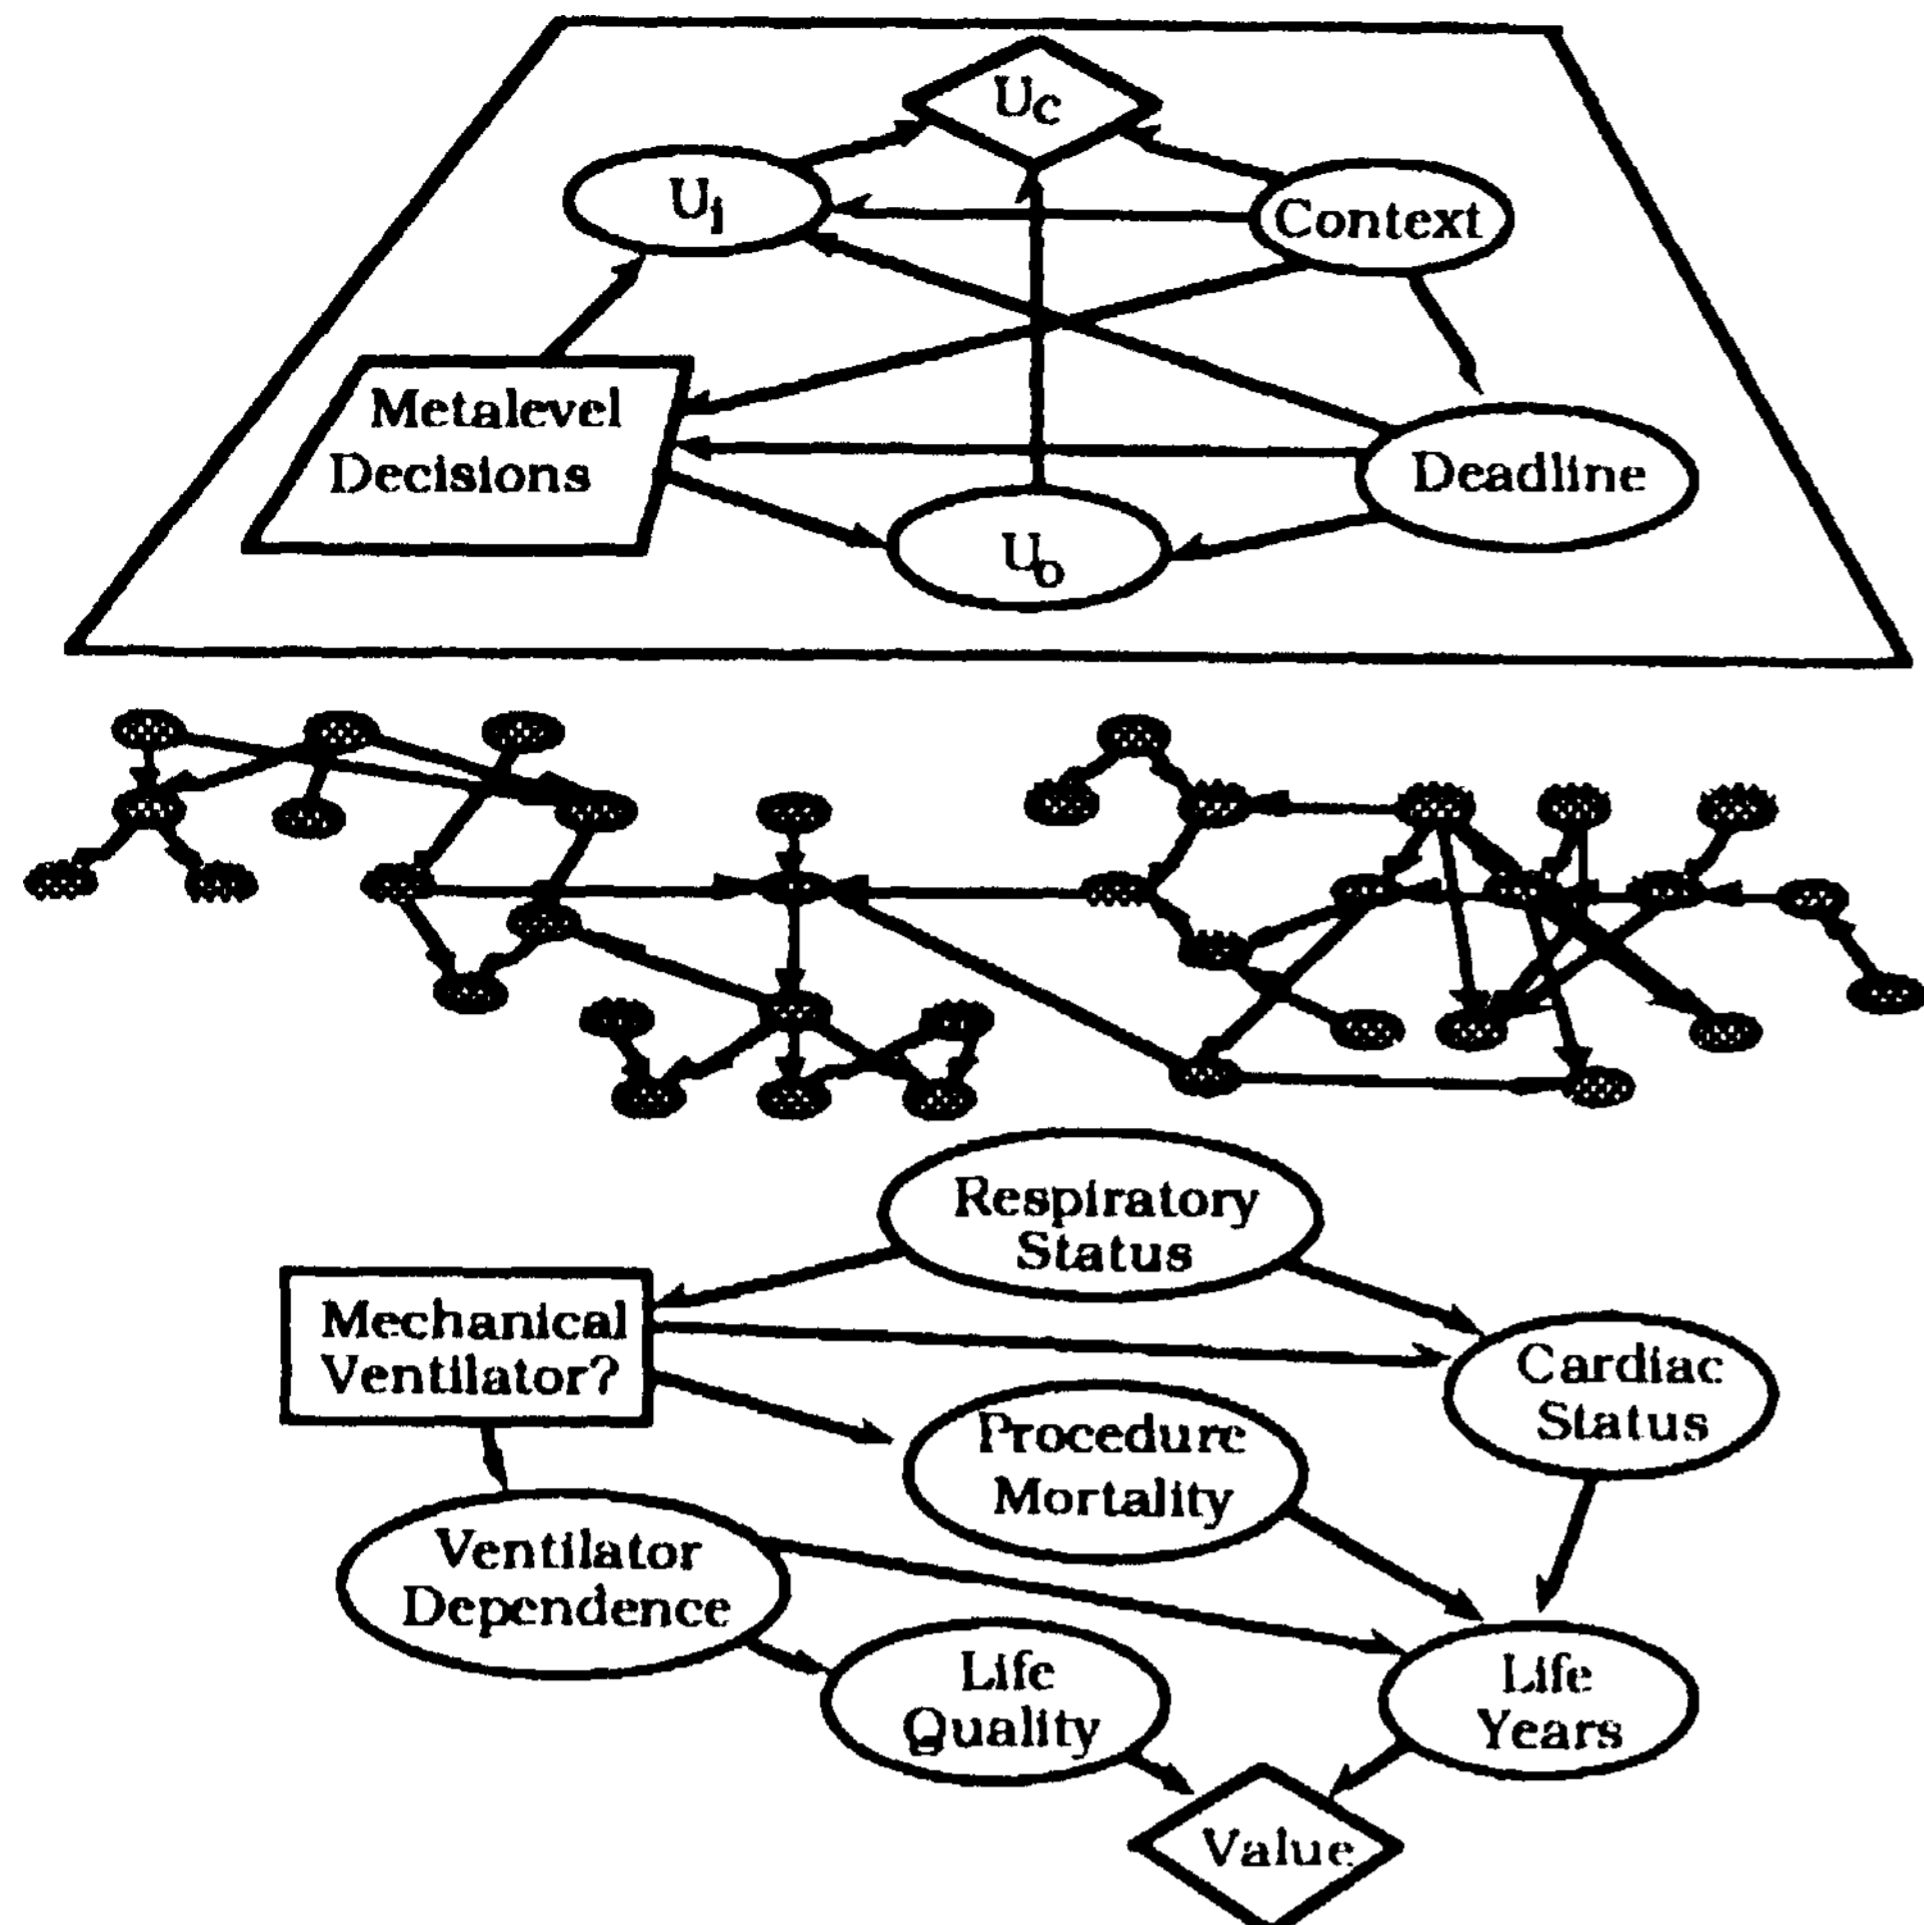
Figure 1: A representation of a time-pressured decision problem. From top to bottom, the three sections of the figure portray (a) the decision-theoretic metareasoning problem, (b) a belief network representing propositions and dependencies in intensive-care physiology, and (c) a closeup on the *respiratory status* node, and its relationship to the current decision problem.

portion of Figure 1 depicts an object-level influence-diagram representation of a time-pressured problem that might face an automated physician's assistant: A 75-year-old woman in the intensive-care unit suddenly shows signs of breathing difficulty. The patient may be merely showing signs of mild *respiratory distress* or may be in the more serious situation of *respiratory failure.* In this context, the primary decision is whether or not to recommend that the patient be placed on a mechanical ventilator. The decision (square node) depends on the probability of respiratory failure, which, in turn, depends on the probabilities of propositions in a large belief network serving as a medical knowledge base (represented by the graph above the object-level problem in Figure 1). The large oval nodes in the base decision problem represent uncertain states associated with placing an older person on a ventilator. The diamond represents the utility associated with different outcomes. Factors to consider in a decision to act include the possibility that it may take a long time to wean a patient with severe lung disease from a ventilator that is applied needlessly; thus, the patient may face a long hospital stay and be placed at high risk of mortality from a disease such as pneumonia. However, if the patient turns out to be in respiratory failure, and is not treated immediately, she faces a high risk of cardiac arrest based on the disrupted physiology associated with abnormal blood levels of oxygen and carbon dioxide.

4.1   Actions and Outcomes in the World

In our simple example, there are only four different fundamental outcomes. The patient either is in respiratory failure *(H)* or is not in respiratory failure $(\neg H)$, and we either will place the patient on a ventilator (.4) or will not do so $(\neg A)$. Thus, we may erroneously decide not to treat a patient who is suffering from respiratory failure $(\neg A, H)$, we may correctly treat a patient who is suffering from respiratory failure *(A,H)*, we may erroneously treat a patient who is not suffering from respiratory failure $(A, \neg H)$, or we may correctly forego treating a patient who is not suffering from respiratory failure $(\neg A, \neg H)$. The expected object-level utilities of action $[u(A)]$ and of no action $[u(\neg A)]$ in terms of the probability of respiratory failure, P(H), are described by the following equations:

$$u(A) = p(H)u(A, H) + p(\neg H)u(A, \neg H)$$

$$u(\neg A) = p(H)u(\neg A, H) + p(\neg H)u(\neg A, \neg H)$$

The lines described by these equations intersect at a probability of *11* denoted *p**. The desired action (the decision with the highest expected utility) changes as the utility lines cross at ;P*. A utility analysis dictates that a patient should not be treated unless a decision maker's belief in the truth of *H* is greater than *p**.

4.2   Decisions About Computation

Let us now integrate explicit knowledge about the process of reasoning into the decision problem. In answer to a query for assistance, our automated reasoner must propagate observed evidence about the patient's symptomology through a complex belief network. The results of an approximate probabilistic-inference scheme may be a probability distribution over a *final* probability. This probability is the value that a computer will calculate from a belief network, given sufficient time to finish its computation. Assume that our reasoner may apply one of several incremental-refinement algorithms that can iteratively tighten the distribution on the probability of interest over time. We wish the system to make a rational decision about whether to make a treatment recommendation immediately, or to defer its recommendation and continue to reason, given its knowledge about the costs of time needed for computation.

4.2.1   Costs of Inference-Based Delay

The example of a patient gasping for breath, facing the risk of a long hospitalization or a cardiac arrest depending on our decision, poignantly demonstrates the salience of reasoning-resource constraints in a high-stakes situation. Sc far, we have considered the utilities of alternative outcomes to be independent of time. Assume that the utility of treating a patient in respiratory failure depends on how long the patient has been in failure. Assume, also, that the initial presentation of respiratory symptoms occurs in the presence of the reasoner and that analysis of the problem begins at this time, t_o. We represent the cost of delaying treatment, when that treatment is needed, by considering a continuum of mutually exclusive decisions to treat at different times, A(i).   where $t = t_o + \Delta t$. A cost function can capture

the decay of utility of action with time. At some time t, the utility of acting in the presence of respiratory failure reverts to the utility of not acting at all. We substitute the static utility equation for *u(A)*, defined previously, with a time-dependant equation:

$$u[A(t)] = p(H)u[A(i),H]+p(-,H)u(A,->H)$$

where u[A(t), *H]* reverts to *u(¬A, H)* as some function of time[2]. In this example, we assume that delay of action will not affect the utility of a patient that does not require the intervention. With the time-dependent utility function, our p* threshold will change with time.

As indicated by the network in the upper portion Figure 1, a more complete representation of the respiratory decision problem includes knowledge about the costs and benefits of applying different inference strategies. This influence diagram represents the metareasoning problem. The node labeled U_o in the metareasoning network is just the value node from the object-level decision problem represented at the bottom of Figure 1. Rather than seek to optimize the object-level value, our agent's goal is to optimize the utility associated with the value node in the metareasoning problem, labeled U_c. As demonstrated by the relationships among propositions in the metareasoning problem, U_c is a function of the object-level value and the inference-related cost, Ui, which in turn depends on computational delay, time availability, and the context. The integration of inference-related and object-level utility allows agents to treat decisions and outcomes regarding the control of reasoning just as it does decisions about action in the world.

4.2.2 Reflection About Future Belief

Our agent's attention is centered on the calculation of $p(H)$, the *probability of respiratory failure.* We define $\lessdot\phi$ to be the probability that the agent would compute if it had sufficient time to finish its computation. That is, ϕ is value of *p(H)* that the reasoner will report after complete computation. At the present moment—before the inference is completed—our automated reasoner may be uncertain about what the value of ϕ will be. The current uncertainty is described by some probability distribution over ϕ. We denote the uncertainty about ϕ *at the present moment* by $p(\phi)$. Although this distribution can change with reasoning, investigators [Howard, 1970] have shown that the belief a decision maker should use for decision making, if she has to act immediately, is the mean of $p()$, denoted by $< \phi >$. After spending additional time *t* on inference about ϕ, our reasoner may have a new distribution over ϕ, denoted by $p_t(\phi)$.

An automated reasoner may have useful knowledge about how a distribution over a belief—and thus how the new mean of the distribution -will change with additional computing. An important class of knowledge about ϕ is of the form, $p(p_t(\phi))$. This measure refers to belief *at the present time* about the likelihood of alternative belief distributions over ϕ that might be generated

In this domain, we could capture the cost of delay with a stochastic model describing the probability of a cardiac arrest as a function of the time we delay therapy; cost models can be useful summaries of the utility of a large number of outcomes.

after computation for additional time *t*. This notion is central in reflection about the value of initiating or continuing decision-theoretic inference, as opposed to that of acting with the current best decision.

Expected Value of Perfect Computation Suppose that, after thinking for only a few milliseconds, an automated reasoner has generated a probability distribution over ϕ. We first introduce the *expected value of perfect computation* on ϕ, denoted by $EVPC_\phi$. The $\mathbf{EVPC}_\phi$ may be viewed as the value of instantaneous complete computation of the target probability in a decision setting. Instantaneous complete thinking would collapse the current probability distribution over ϕ into an impulse. Given the current probability distribution $p(\phi)$ over ϕ, we define $EVPC_\phi$ as follows:

$$\mathbf{EVPC}_\phi = \int_\phi p(\phi) \max_D u_o[D(\phi)]\, d\phi - \max_D u_o[D < \phi_o >]$$

where $\max_D u_c[D < \phi_o >]$ is the utility, associated with the best decision *D,* based on taking an immediate action using the current mean belief, $< \phi_o >$. This measure tells us that the value of computing the final answer is just the difference in utility between the current best action and the summation of future best actions weighted by the probability of different final beliefs.

Belief About Changes in Belief Real-world computers rarely deliver the full expected value of perfect computation on difficult problems because they must expend valuable resources in the reasoning process. Assume that our agent in the intensive-care unit, faced with determining the probability that our elderly patient is in respiratory failure, has incomplete knowledge about what $p(\phi)$ will be at some future time t, which we refer to as $p(p_t(\phi))$. For example, the system may have a probability distribution over the future bounds on ø with additional computation. Such knowledge may have been acquired through an empirical analysis of a network in addition to an upper bound that has been proved theoretically. Our reasoner could apply this type of knowledge by considering the EVC(t) based on the information about probability distributions over $p(\phi)$, obtained with computation for an additional time *t,* as

$$\mathbf{EVC}(t) = \int_{p_t(\phi)} p(p_t(\phi)) \int_\phi p(\phi) \max_D u_c[D(\phi),t]\, d\phi\, dp_t(\phi)$$

$$- \max_D u_o[D(< \phi_o >)]$$

That is, we sum over the new probability distributions on ϕ expected at time *t,* weighted by the *current* belief, $p(p_t(\phi))$, that thinking until *t* will lead to each of the revised distributions. In terms of the mean, $< \phi_t >$ of the future distributions, $p_t(\phi)$,

$$\mathbf{EVC} = \int_{p_t(\phi)} p(p_t(\phi)) \max_D u_c[D(< \phi_t >,t)]\, dp_t(\phi)$$

$$- \max_D u_o[D(< \phi_o >)]$$

When, for all *t,* the cost of computation, embodied within our comprehensive utility function, becomes
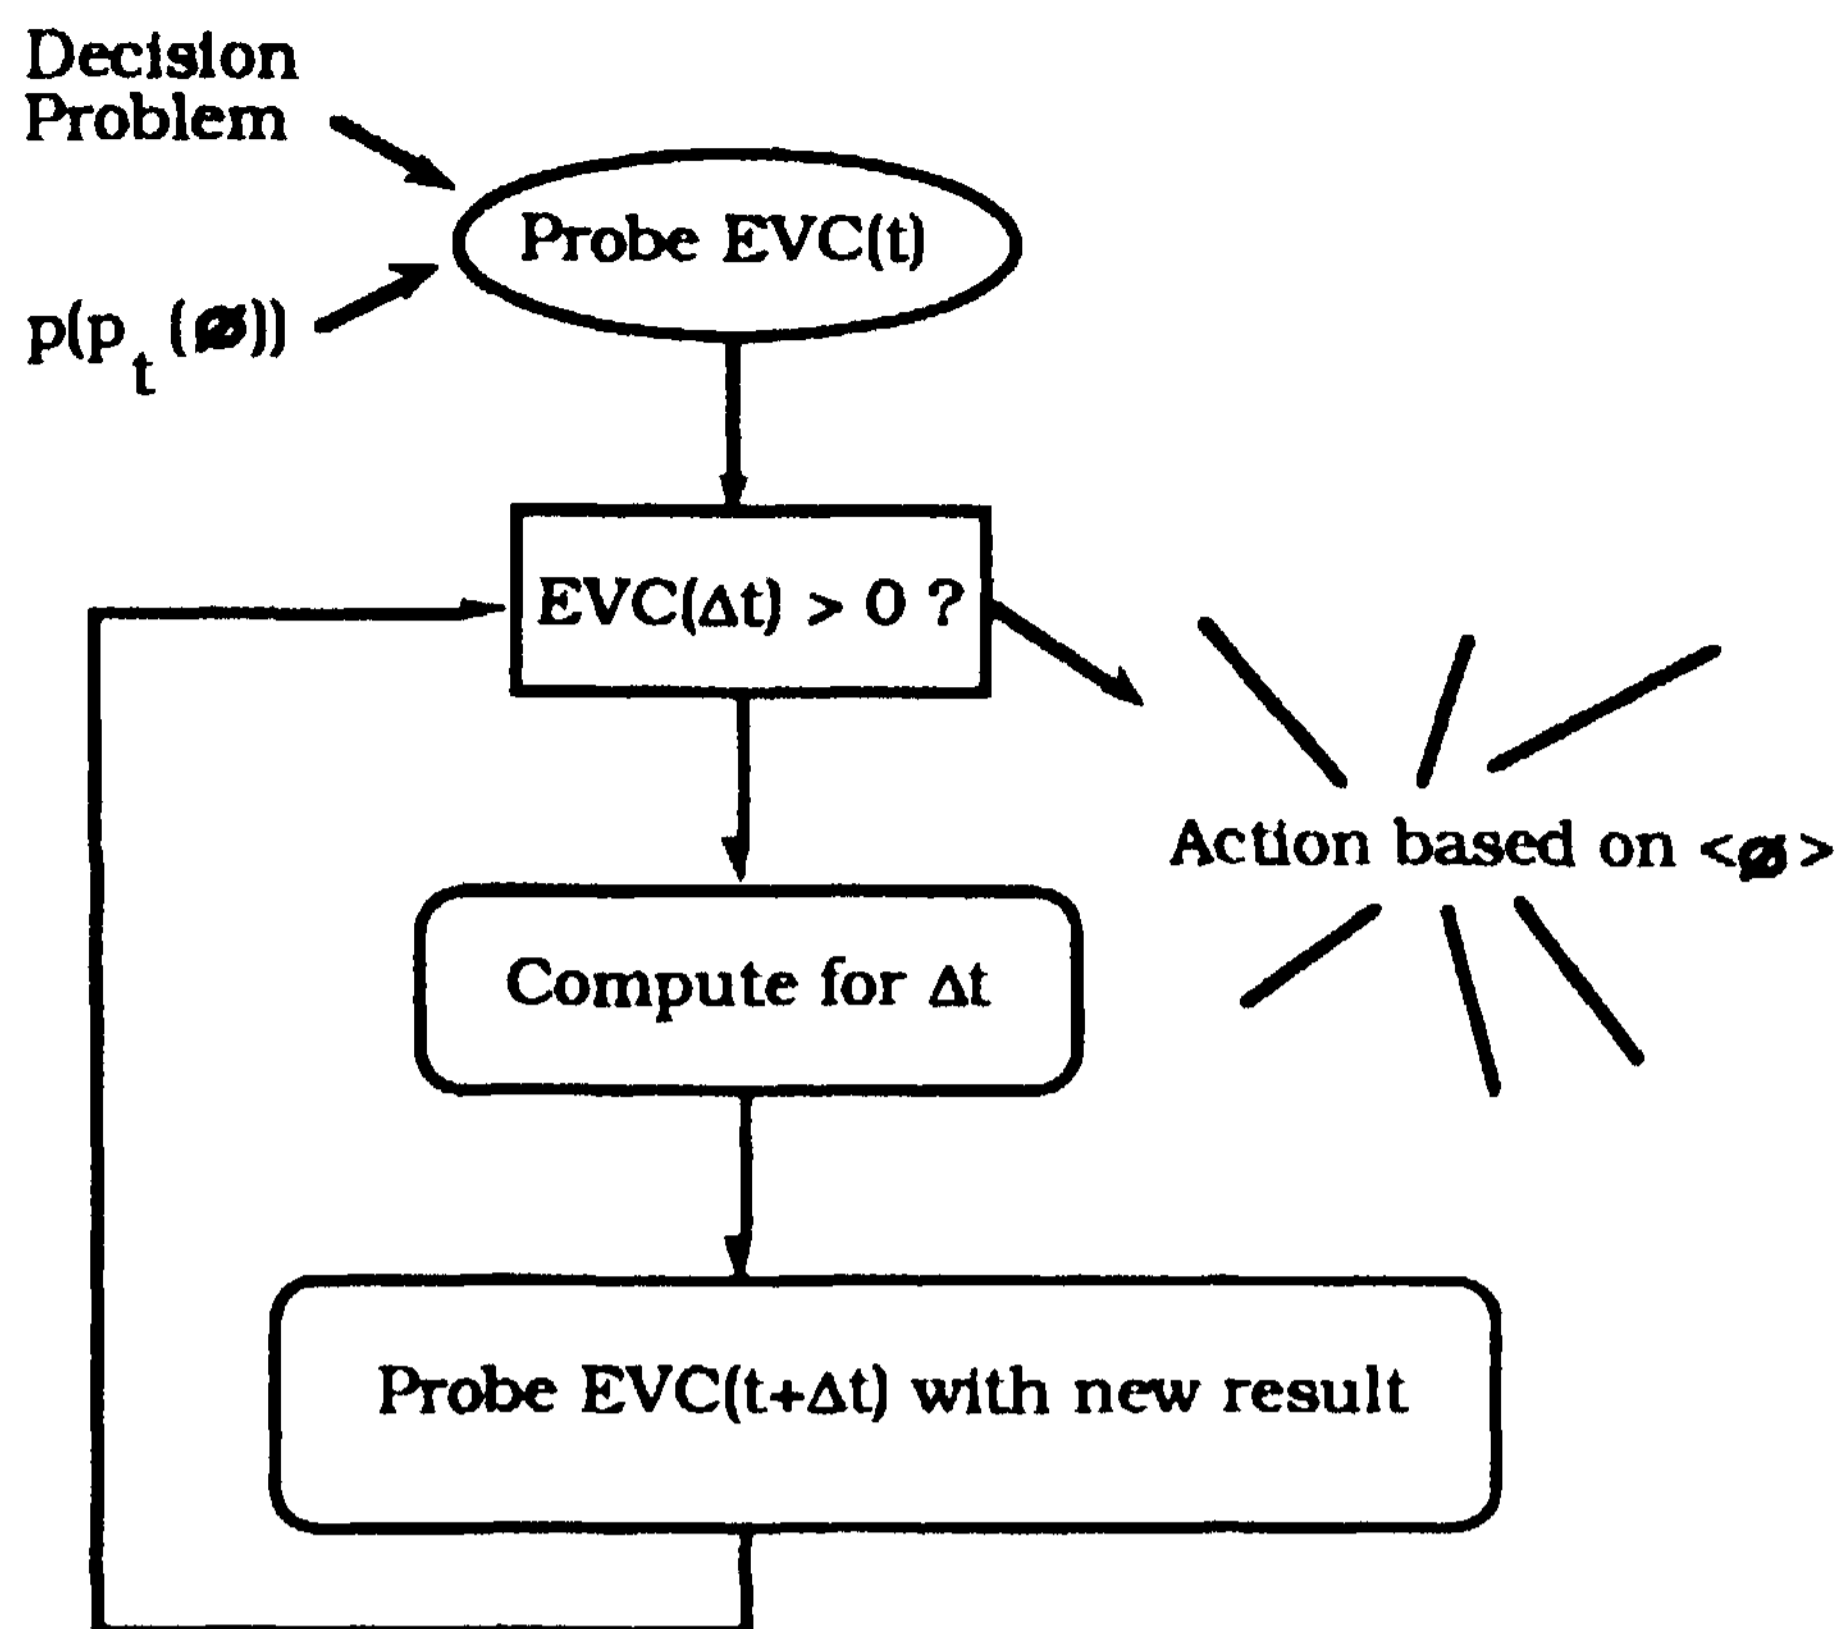
Figure 2: In reasoning about the value of continuing to reflect about belief, versus that of taking immediate action, an EVC-evaluation module considers the decision problem and the probability distribution over future probability distributions [p(pt (ø))] that may be generated with the allocation of computational resource.

greater than the benefit of computing (EVC < O), an agent should cease reflection and act. The EVC formula can be used to study the value of alternative inference schemes. Of course, there can be significant overhead in the real-time application of an EVC-based control strategy. Thus, a central goal of research on decision-theoretic control is to identify tractable solutions to the EVC evaluation problem. Alternatively, offline analysis and compilation of control strategies may be useful in situations where the complexity of meta-analysis limits the gains of real-time decision-theoretic control.

Analogous value-of-computation approaches can be used to valuate and control other problem classes. For example, we can use an *EVC(t)* calculation for controlling the nature and extent of a search or sort problem; we associate a cost with the time required to expand another node in a tree, or to perform a set of tests and swaps in a partially sorted file, and consider a probability distribution over the expected object-level gains, given the allocated time. The development of tractable EVC approximations for these and other problems, make possible useful normative control through iterative testing of the value of continuing to reason.

5   Value of Probabilistic Bounding

We have pursued tractable solutions to the EVC through examining parameterized families of distributions. For example, we have explored the use of the EVC approach to control probabilistic bounding methods. Assume that our automated reasoner has, with some initial amount of computation, computed upper and lower bounds on ø), with an upper bound at *b* and lower bound at a. If our

reasoner does not have any information about where ø is—except that the final computed result will be between the current bounds—then it is reasonable to assume a uniform distribution over ø between the bounds. A uniform distribution within bounds is consistent with an agent being ignorant of the final belief, except for the bounds information. Detailed knowledge about convergence could change this distribution. Let us focus on the value structure of assuming uniformity at both current and future distributions about belief. We denote the problem by EVC/BU, the expected value of computation for a bounding algorithm given an assumption of uniformity.

An agent may have useful knowledge about $p_t(ø)$ *without* having information about how the mean $< ø_t >$ will change, except for knowing that ø will be constrained to tighter bounds. As an example, a system could make use of certain or uncertain knowledge about the *rate* of bounds convergence to valuate a decision to continue to compute. We have analyzed how a system can apply knowledge that the bounds on the belief for a node in a belief network will converge at a rate dictated by a fraction, C, which, when multiplied by the current bounds interval at t_o, dictates the interval at *t*. That is,

$$int_t = C * int_{t_o}$$

where *int* is just the interval, or the difference between the upper and lower bounds. If we were uncertain about the convergence, we would have a probability distribution over this convergence fraction.

We have applied the EVC equation to the bounding problem, considering future distributions expected with additional computation. Given a convergence fraction that allows us to calculate the future bounds, we must consider all possible configurations of the new bounds given the current constraints. As we sweep the expected future interval over the current interval, the mean of the future distribution sweeps between positions within the current bounds. When the mean is above p*, we sum over the utility of acting for all states of belief greater than that threshold; when the mean is below *p**, we similarly consider the utility of not acting. Given our current bounds and a convergence fraction, we sum the utilities of the best decision at the future means and subtract the utility of the best action without additional computation. Solving the uniform distribution case for different possible p* boundary conditions yields functions that report the EVC as a function of (1) the utilities for each of the four outcomes, (2) the current bounds on ø, (3) a function describing the expected convergence of bounds (e.g., a C, with time, and (4) the cost of delay. Under uncertain performance, a rational agent's reflection based on the EVC formalism involves the interlacing of probes for positive EVC(t) and continued inference. Computation should continue until action is indicated by a non-positive EVC. This volley of reflection and inference is demonstrated in Figure 2.

5.1   Partial Characterization of Inference

We have experimented with decisions about computation and action within alternative utility contexts. We have

particularly explored the behavior of recently-developed graceful approximation methods for probabilistic inference. These strategies include a flexible variant of Pearl's method of conditioning [Pearl, 1986], called *bounded conditioning* [Horvitz et al., 1988b].

In the method of conditioning, a multiply connected network is reformulated to a set of singly connected networks by locating a set of nodes that break cycles. The complete set of cycle-breaking nodes is called the *loop cutset*. The nodes of the loop cutset are instantiated with each possible value (or combination of values), and the resulting joint probabilities of each *instance* are calculated as prior probabilities of the instantiated variables. Algorithms for solving the singly connected network subproblems can be applied to the solution of each network instance. In bounded conditioning, instances are analyzed in order of their expected contribution to the tightening of bounds. The instances are sorted according to their prior probability, and are solved in sequence. A bounding calculus generates logical bounds on the final probability of interest by considering the maximum and minimum contributions of the unexplored subproblems.

We applied bounded conditioning to several random networks as well as to a belief network describing probabilistic relationships among findings and pathophyisologic states in an intensive-care unit.[3] The structure of this belief network is captured by the graph in the middle of Figure 1. The network consists of 37 multiply connected nodes. We studied the performance of several loop cutsets for this network. A sample loop cutset consists of 5 nodes that leads to 144 different singly connected-network problems.

We sought to characterize the refinement of bounds with additional computation. Our analyses focused on updating belief in the intensive-care network with single pieces of evidence. We found that the convergence of the bounds could by approximated by an exponential decay of the size of the interval with time. This convergence was modeled approximately by the function

$$int = e^{-k(t+1)}$$

Additional discussion of bounded conditioning, including analysis of the basis for such convergence, is found in [Horvitz et al., 1988b]. As an example, the convergence of a typical update in the network is captured by an exponential decay with an approximate half-life of 36 seconds. That is, after 36 seconds of analysis by a Motorola-68020-based computer, running at a 17 MHz clock rate, the bounds converge to one-half of their original bounds. At 72 seconds, the bounds are halved once again to an interval of approximately 0.25. This convergence is modeled by the exponential decay with $k = 0.02$. The convergence is displayed in Figure 3.

This convergence information can be used to calculate an EVC associated with continuing to apply the bounding algorithm. Evaluating the EVC within our testbed intensive-care belief network has shown, for sample updates and associated sets of utility estimates, that a p* decision threshold can be crossed well before the

[3]This network, called ALARM, was constructed by Ingo Beinlich [Beinlich et al., 1989].

Figure 3: The application of the bounded-conditioning method to the intensive-care unit belief network problem. The graph shows the convergence of the upper and lower bounds, and the mean of the approximation (center curve), to a probability of interest as additional inference subproblems are solved.

computation of final belief. Experimentation with rational metareasoning to select among alternative inference strategies, and to control the length of time that they are applied is continuing on a variety of belief networks and decision contexts.

5.2 Acquisition of Control Knowledge

Our formalism for the calculating the value of probabilistic bounding operates on knowledge about convergence on bounds. We have performed theoretical analysis of worst-case performance of bounded conditioning. We have also recorded empirically derived partial-characterization information. Clearly, an agent could benefit by continually bolstering its knowledge about partial characterizations with extensive empirical study of problem-solving trajectories during idle-time. A component of our research focuses on an offline analysis of the performance of reasoning strategies of different networks. The analyses are aimed at capturing useful partial characterization of the expected performance of different strategies by performing Monte Carlo simulation to generate plausible patterns of evidence, and summarizing and storing a set of performance indices. For example, we are interested in the convergence of bounds in response to a state of evidence. This information can be extremely useful to a control reasoner that is attempting to valuate the EVC for a set of competing solution strategies. We are researching the automated acquisition of partial characterizations of stategy performance within the intensive-care unit application area.

6 Summary

We have described research on the rational interlacing of decision-theoretic inference with action under scarce resources. We highlighted the use of partial characterizations of probabilistic inference to reason about the value of continuing to reason about a problem versus that of taking action in, the world. After defining the expected value of perfect computation, we explored the expected value of computation for reasoning about belief in a decision context. We focused on the valuation of future computation, based on a consideration of future probability distributions over the truth of a proposition of interest. We described the use of normative metareasoning techniques for valuating and controlling probabilistic inference for time-pressured medical decisions. Preliminary empirical work has demonstrated that a probability bounding algorithm can deliver a large fraction of the expected value of perfect computation well in advance of complete inference. We are continuing to experiment with different distributions over belief and are working to characterize useful dimensions of algorithm performance. We foresee that advances in the application of decision-theoretic reflection will play an important, role in the development of effective normative reasoning systems.

7 Acknowledgments

We thank Ronald Howard, Edward Shortliffe, and Jacques Suermondt for their feedback on this work. Lyn Dupre provided comments on an earlier draft.

References

[Beinlich et al., 1989] I.A. Beinlich, H.J. Suermondt, R.M. Chavez, and Gregory F. Cooper. The ALARM monitoring system: A case study with two probabilistic inference techniques for belief networks. Proceedings of AI in Medicine, London (in press).

[Cooper, 1987] G.F. Cooper. Probabilistic inference using belief networks is NP-hard. Technical Report KSL-87-27, Knowledge Systems Laboratory, Stanford University, Stanford, CA, May 1987. Also in Artificial Intelligence (in press).

[Dean and Boddy, 1988] T. Dean and M. Boddy. An analysis of time-dependent planning. In Proceedings AAAI-88 Seventh National Conference on Artificial Intelligence, pages 49 54. American Association for Artificial Intelligence, Minneapolis, MN, August 1988.

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

[Good, 1968] I.J. Good. A five-year plan for automatic chess. In Machine Intelligence, 2:89 118, London, Oliver and Boyd, 1968.

[Hansson and Mayer, 1989] O. Hansson and A. Mayer. Probabilistic heuristic estimates. In Proceedings of the Second International Workshop on AI and Statistics, Ft. Lauderdale, January 1989.

[Horvitz et al., 1988a] E.J. Horvitz, J.S. Breese, and M. Henrion. Decision theory in expert systems and ar-tificial intelligence. International Journal of Approximate Reasoning, 2:247-302, 1988.

[Horvitz et al., 1988b] E.J. Horvitz, H.J. S uermondt, and G.F. Cooper. Bounded conditioning: Flexible inference for decisions under scarce resources. Technical report KSL-88-36, Stanford University, July 1988.

[Horvitz, 1987a] E.J. Horvitz. Problem-solving design: Reasoning about computational value, tradeoffs, and resources. In Proceedings of the NASA Artificial Intelligence Forum, pages 26-43. National Aeronautics and Space Administration, Palo Alto, CA, November 1987.

[Horvitz, 1987b] E.J. Horvitz. Reasoning about beliefs and actions under computational resource constraints. Third Workshop on Uncertainty in Artificial Intelligence, Seattle, WA, July 1987. American Association for Artificial Intelligence. Also in L. Kanal, T. Levitt, and J. Lemrner, ed., Uncertainty in Artificial Intellipfnee 3, Elsevier, Amsterdam (in press).

[Horvitz, 1988] E.J. Horvitz. Reasoning under varying and uncertain resource constraints. In Proceedings AAAI-88 Seventh National Con ferenee on Artificial Intelligence, pages 111 116. American Association for Artificial Intelligence, Minneapolis, MN, August, 1988.

[Howard and Matheson, 1981] R.A. Howard and J.E. Matheson. Influence diagrams. In R.A. Howard and J.E. Matheson, editors, Readings on the Principles and Applications of Decision Analysis, volume II, pages 721-762. Strategic Decisions Group, Menlo Park, CA, 1981.

[Howard, 1970] R.A. Howard. Decision analysis: Perspectives on inference, decision, and experimentation. Proceedings of the IEEE, 58:632-643, 1970.

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

[Pearl, 1986] J. Pearl. A constraint-propogation approach to probabilistic reasoning. In L.N. Kanal and J.E. Lemmer, editors, Uncertainty in Artificial Intcl-11gence I, pages 357 369. Elsevier, Amsterdanl, 1986.

[Russell and Wefald, 1988] S.J. Russell and E. Wefald. Multi-level decision-theoretic search. In AAAI Spring Syrnpcsium on Came Playing, Stanford, CA, March 1988. American Association for Artificial Intelligence.

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

[Treitel and Genesereth, 1986] R. Treitel and M.R. Genesereth. Choosing directions for rules. In Proceedings AAAI-86 Fifth National Conference on Artificial Intelligence, Philadelphia, PA, August 1986. American Association for Artificial Intelligence.

[Watson and Brown, 1978]
S.R. Watson and R.V. Brown. The valuation of decision analysis. J.R. Statist. Society A., 141:69-78, 1978.