Research on Principles of Bounded Rationality

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

The Complexity of Normativity

Several years ago, we initiated the PROTOS project, to investigate the use of decision theory as a framework for reasoning about the design and operation of ideal expert systems and control systems under bounded resources. The research evolved from early research on the PATHFINDER project, an effort to development an expert system for surgical pathology. PATHFINDER investigators have investigated the use of probability theory for representing and manipulating *uncertain* knowledge. Intuitive and theoretical justification of the axioms of probability and utility theory have led people in several disciplines to view the axioms of decision theory as *normative*.

A chief problem with the development of normative reasoning systems is the *complexity* of traditional normative representations and inference methods. Several simplified normative medical diagnostic systems were developed during the 1960s and early 1970s. Attempts to extend the scope of these systems, or to weaken the simplifying assumptions of manifestation independence, and mutual exclusivity of diseases, threatened researchers in artificial intelligence in medicine (AIM) with combinatorial explosion in the effort required for representation and computation.

In PATHFINDER and related PROTOS research, we have wrestled with the problems of complexity *within* the framework of probability and decision theory. We have addressed complexity in two different ways. A large portion of PATHFINDER research has grappled with complexity by developing representations and tools that increase the efficiency of elucidating and capturing *independence* in the structure of biomedical knowledge [Heckerman *et al.*, 1989]. That work has demonstrated, for several domains, that there is a great deal of independence in expert knowledge. Unfortunately, we may not always be able to rely on independence to gain tractability; medical knowledge can be intrinsically complex. PROTOS research has investigated the normative control of representation and inference problems that are too complex to solve without incurring unacceptable computational or engineering costs. This research has sought to develop principles of bounded rationality based on decision theory, for use in controlling tradeoffs associated with the use of approximate normative reasoning strategies.

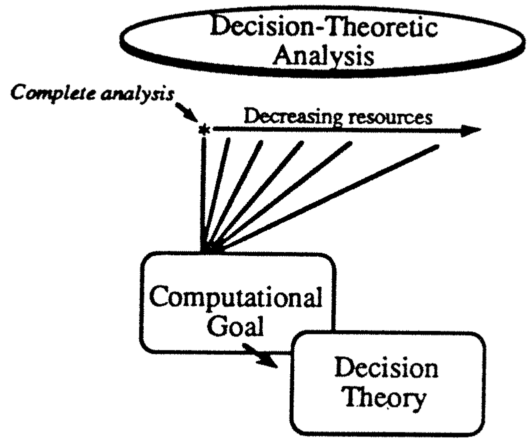

Figure 1: Rational metareasoning about partial computation can be applied to such fundamental operations as sorting and searching, or to decision-theoretic inference itself. We wish to optimize the value of an approximation, as the solution ranges from a complete analysis (asterisk) to an increasingly poorer approximation (direction of arrow)

Bounded Optimality as Rationality

The high stakes and time pressure of decisions in medicine highlight the significance of developing nor-

*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 NLM Grant LM05208.

mative principles of bounded rationality. As indicated in Figure 1, our model of bounded rationality centers on the use of design-time and tractable run-time decision-theoretic analyses to control the detail and completeness of complex problem-level decision making.

This normative perspective on bounded rationality is in sharp distinction to older, more familiar conceptions of bounded rationality. In 1955, Herbert Simon noted that we should consider constraints on cognitive resources in generating and evaluating the behavior of a decision maker in a complex situation. However, he and other early AI pioneers quickly retreated from decision theory. Citing the limited abilities of human decision makers and the forgiving nature of many problems in the world, Simon proposed that most intelligent behavior is oriented toward finding relatively simple solutions that are nonoptimal, yet are sufficient or *satisficing* [Simon, 1955]. This theme has stimulated broad AI research on relatively ill-characterized heuristic procedures in a wide array of domains. Unfortunately, the nonnormative approaches to *bounded rationality* 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, given complex uncertainties about the world. In medical reasoning, potential losses—and opportunities for great gain—highlight the potential usefulness of decision theory for optimizing the value of behavior under resource constraints.

We use *bounded optimality* to distinguish research on the optimal design of problem solvers and solution methodologies under bounded resources from traditional nonnormative approaches to reasoning under resource constraints [Horvitz, 1987]. Unlike straightforward decision analyses, we apply the principles of decision theory to enriched models that include not only distinctions and outcomes in the world, but also distinctions and outcomes about the representations or reasoning processes themselves.

Focus: Inference in Influence Diagrams

We have been investigating the use of *influence diagrams* for representing and solving difficult medical reasoning problems. The influence diagram is an acyclic directed graph containing nodes representing propositions and arcs representing interactions 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. 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.*, 1988]. 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 flexible approximation strategies and intelligent control techniques.

Utility of Computation

As background, we will briefly describe fundamental concepts of decision-theoretic control and flexible computation. We use *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_o, 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. The *inference-related* component is 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.

Flexible Inference and Representation

A important aspect of developing reasoners that are resilient to uncertain challenges and resources is the development of *flexible* computation and representation strategies.

Desiderata of Flexible Computation

We have enumerated and analyzed properties of flexible reasoning that are desirable for reasoning under bounded resources [Horvitz, 1987,Horvitz, 1988]. First, we wish to see *monotonicity* in the change of the object-level value of results with the expenditure of resource. We desire our strategies also to exhibit graceful degradation, or to be relatively insensitive to small reductions in the allocated resource. That is, we prefer incrementality or *continuity* in the refinement of partial results with the application of resources. We refer to strategies that exhibit *monotonicity* and *continuity* with computation as *flexible-computation* policies.[2] *Spanning* strategies are

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

[2] Dean and associates later independently referred to monotonic-refinement policies, in the context of planning

an especially valuable class of flexible-computation policies that converge to an ideal solution with some finite allocation of resource.

Economics of Flexible Computation

As an example of some prototypical relationships, consider an example of the costs and benefits of reasoning, borrowed from [Horvitz, 1987]. Figure 2 highlights the fundamental relationships among u_o, u_i, and u_c. As indicated by the figure, the ideal halting time, t^*, for problem-solving methods, described by a comprehensive utility function, with an object-level component that is monotonically increasing and a negative second derivative, and an inference-related utility component that is monotonically increasing, with a second derivative that is zero or positive, occurs when the first derivatives of both components are equal. We partition the t into with computation time, t_c, and metalevel reasoning time, t_m. Consider the case where the object-level value is refined with computation time t_c as a negative exponential process

$$u_o = 1 - e^{-kt_c} \tag{1}$$

Let us assume that the cost of delay is separable from the object-level utility and that the object-level value and cost are related by addition. Assuming that we have a linear cost with

$$u_i = -ct \tag{2}$$

we know that the comprehensive utility is

$$u_c = 1 - e^{-k(T-t_m)} - ct \tag{3}$$

where t_m is the constant cost of the metalevel optimization process needed to determine the optimal halting point. We can solve for an optimal halting time by differentiating Equation 3 in terms of the time expended on the problem and identifying a maxima. For our example, the ideal amount of computation time before halting is

$$t_c^* = \frac{-ln(\frac{c}{k})}{k} \tag{4}$$

Through substitution, into Equation 3, we can calculate the ideal u_c, u_c^*, at this ideal halting point. This optimal utility is

$$u_c^* = 1 - c(\frac{1 - ln(\frac{c}{k})}{k} + t_m) \tag{5}$$

These equations demonstrate how the optimal halting time and ideal comprehensive utility will change, given changes in the cost of reasoning (dictated by c), in the rate of refinement (dictated by k), and the cost of determining the optimal halting time (t_m).

problems, as *anytime algorithms* [Dean and Boddy, 1988].

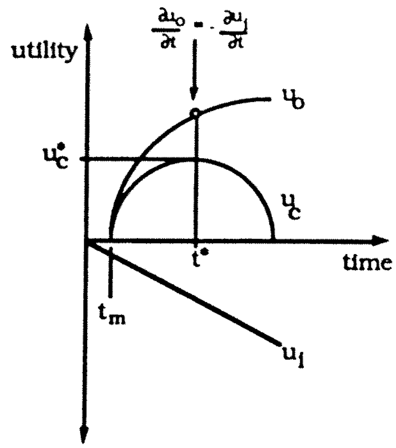

Figure 2: The economics of flexible computation. The graph highlights the fundamental relationships among the object-level value (u_o), the cost of reasoning (u_i), and the comprehensive value of computation (u_c).

Flexible Probabilistic Inference

The previous example assumes certainty in the object-level value and costs. We are interested more typically in reasoning under uncertainty. Our goal has been to develop flexible methods for probabilistic inference and to develop theoretical and practical tools for the control of the flexible reasoning. As an example, we have constructed an algorithm for flexible probabilistic inference named *bounded conditioning* [Horvitz *et al.*, 1989c]. Bounded conditioning satisfies the desiderata of *continuity*, *monotonicity*, and *convergence*. The method works by decomposing an inference problem into separate subproblems, each representing a plausible context, and generating exact bounds on probabilities of interest through accounting for the contexts that have not yet been explored. The algorithm incrementally refines bounds on a probability of interest, continuing to tighten the upper and lower bounds on a probability of interest, with continuing computation, until reaching a point probability.

We characterized the behavior of bounded conditioning on experimental belief networks constructed with a belief-network generator and on an intensive-care unit (ICU) belief network [3] Figure 3 shows the typical form of the convergence displayed by bounded conditioning when updating all nodes in the ICU network, given an observation. We found that the bounds for a large set of updates for this problem decay at rate that can be modeled approximately with a negative exponential, e^{-kt}, with different decay constants k.

[3]The network was developed by Ingo Beinlich, a doctoral student in our research group.

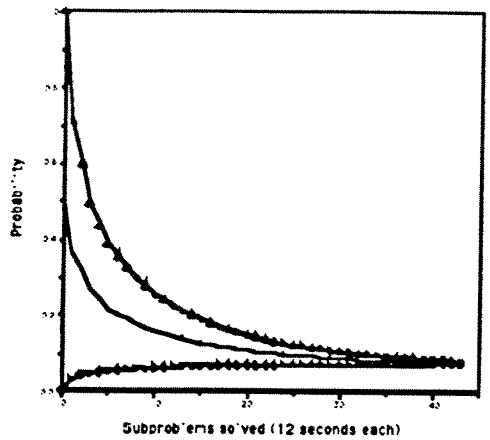

Figure 3: The flexible bounded-conditioning inference strategy incrementally refines bounds on probabilities in a belief network.

Ideal Reflection Before Action

A component of PROTOS work on the metalevel control of probabilistic inference is captured by the overview in Figure 4. A metareasoning decision problem, shown at the top of the figure, is used to control the nature and extent of inference in a complex belief network. The structure of the ICU belief network that we have analyzed is represented in the middle of Figure 4. This network represents uncertain relationships between observations and patient pathophysiology in intensive-care medicine. Object-level decision problems, requiring belief-network inference for relevant probabilities, are passed to the metareasoner, which determines the optimal dwell time.

In answer to a query for assistance, our automated reasoner must propagate observed evidence about a patient's symptomology through the complex ICU belief network. Rather than seeking simply to optimize the object-level value by doing complete inference, the reflective system'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, u_i. The integration of inference-related and object-level utility allows our system to treat decisions and outcomes regarding the control of reasoning just as it does decisions about action in the world.

Given a characterization of the performance of flexible algorithms as a function of the expenditure of computational resource (such as the convergence of bounded conditioning), we can trade-off the quality of a decision for delay. 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.

The reasoner's attention is centered on the calculation of $p(H)$, the *probability of disease H*. We define ϕ to be the value of $p(H)$ a system would compute if it had sufficient time to finish its computation. At times before the inference is completed, our automated reasoner is uncertain about the value of ϕ. The current uncertainty can be described by a probability distribution over ϕ. We denote the uncertainty about ϕ as $p(\phi)$. Although this distribution can change with reasoning, investigators have shown that the belief a decision maker should use for decision making, if she has to act immediately, is the mean of $p(\phi)$, 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 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 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.

Our reasoning system has incomplete knowledge about what $p(\phi)$ will be at some future time t, which we refer to as $p(p_t(\phi))$. The system typically may have a probability distribution over the future bounds on ϕ with additional computation. Our reasoner can apply this knowledge by considering the EVC(t) based on the information about probability distributions over $p(\phi)$, obtained with computation for additional t, as

$$\mathrm{EVC}(t) = \int_{p_t(\phi)} p(p_t(\phi)) \int_{\phi} p(\phi) \max_D u_c[D(\phi), t]\, d\phi\, dp_t(\phi) - \max_D u_o[D(< \phi_o >)]$$

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, $p_t(\phi)$,

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

When, for all t, the cost of computation becomes greater than the benefit of computing (EVC $\leq O$), an agent should cease reflection and act. The EVC formula can be used to study the value of alternative inference schemes.

There can be significant overhead in the metareasoning time, t_m, required for the application of an EVC-based control strategy. Thus, a central goal of research on decision-theoretic control is to identify tractable

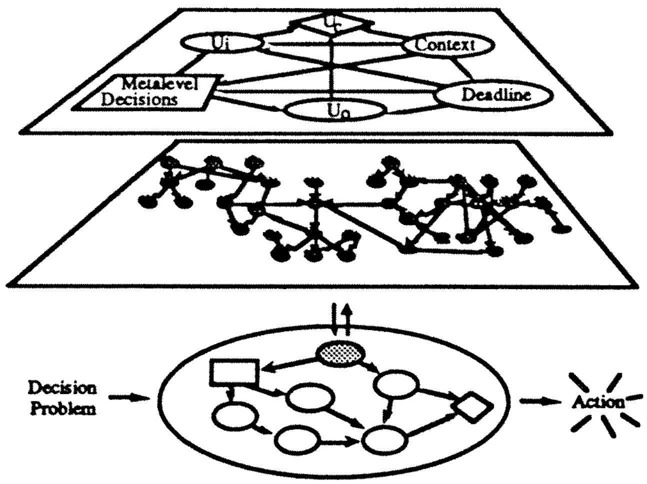

Figure 4: Research on normative metareasoning has explored the use of a metalevel decision model to control probabilistic inference in an object-level belief-network.

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 control. We have pursued tractable solutions to the EVC by examining parameterized families of distributions [Horvitz *et al.*, 1989a]. For example, we have explored the use of rational metareasoning to control the application of probabilistic-bounding methods. The inference metaknowledge is in the form of a partial characterization of how the probability distribution over propositions of interest will change as reflection continues. We found that approximation strategies can typically deliver a large portion of the value of complete computation with a fraction of the resource.

Summary

Theoretical and empirical PROTOS research has demonstrated that principles and models from the decision sciences can be extended and enriched to provide a principled approach to bounded rationality. We have developed techniques to optimize the value of decisions under resource constraints by expanding traditional decision models into more comprehensive analyses. The enriched models include knowledge about problem solving, in addition to domain knowledge. We focused here on the interlacing of flexible computation and normative metareasoning techniques for valuating and controlling probabilistic inference for time-pressured medical decisions. Empirical work has demonstrated that an approximation algorithm can deliver a large fraction of the expected value of perfect computation well in advance of complete inference. Other aspects of research on bounded optimality, with promise for AIM applications, include decision-theoretic techniques for limiting effort in probabilistic assessment, for generating ideal policies for compiling precomputed results, and for trading off the precision for the understandability of explanations of decision-theoretic inference [Horvitz *et al.*, 1989b]. We foresee that advances in the application of decision-theoretic metareasoning will play an important role in the development of effective normative reasoning systems for medicine and other applications dominated by high-stakes decisions under uncertainty.

References

[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. AAAI, August 1988.

[Heckerman *et al.*, 1989] D.E. Heckerman, E.J. Horvitz, and B.N. Nathwani. Update on the pathfinder project. In *Proceedings of the Thirteenth Symposium on Computer Applications in Medical Care*, pages 203–207. IEEE Computer Society Press, November 1989.

[Horvitz *et al.*, 1988] E.J. Horvitz, J.S. Breese, and M. Henrion. Decision theory in expert systems and artificial intelligence. *International Journal of Approximate Reasoning*, 2:247–302, 1988. Special Issue on Uncertain Reasoning.

[Horvitz *et al.*, 1989a] E.J. Horvitz, G.F. Cooper, and D.E. Heckerman. Reflection and action under scarce resources: Theoretical principles and empirical study. In *Proceedings of the Eleventh IJCAI*, pages 1121–1127. AAAI/IJCAI, August 1989.

[Horvitz *et al.*, 1989b] E.J. Horvitz, D.E. Heckerman, K.C. Ng, and B.N. Nathwani. Heuristic abstraction in the decision-theoretic Pathfinder system. In *Proceedings of the Thirteenth Symposium on Computer Applications in Medical Care*, pages 178–182. IEEE Computer Society Press, November 1989.

[Horvitz *et al.*, 1989c] E.J. Horvitz, H.J. Suermondt, and G.F. Cooper. Bounded conditioning: Flexible inference for decisions under scarce resources. In *Proceedings of Fifth Workshop on Uncertainty in Artificial Intelligence*, Windsor, Canada, August 1989.

[Horvitz, 1987] 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.

[Horvitz, 1988] E.J. Horvitz. Reasoning under varying and uncertain resource constraints. In *Proceedings AAAI-88 Seventh National Conference on Artificial Intelligence*, pages 111–116. AAAI, August 1988.

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