Reasoning about Beliefs and Actions Under Computational Resource Constraints[1], pages 312-314.

6.2.2 Bounded Optimality

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

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

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

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

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

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

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

[1]From E. Horvitz, Reasoning about Beliefs and Actions under Computational Resource Constraints. *Third Workshop on Uncertainty in Artificial Intelligence*, Seattle, Washington. July 1987. Association for Uncertainty and Artificial Intelligence. pp 429-444. Also in L. Kanal, et al. eds., *Uncertainty in Artificial Intelligence 3*, Elsevier, 1989, pp. 301-324.

time horizon t, we may prefer the behavior of an agent, A^*, where

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

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

7 Flexible Inference and Intelligent Control

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

...