Qualitative Decision Theory and Interactive Problem Solving (Extended Abstract)

Richmond H. Thomason
Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260
thomason@isp.pitt.edu

Abstract

This paper provides an overview of recent work on qualitative approaches to decision theory. There is a good deal of basic work to be done on the ideas that are emerging from this work before we can apply them with much confidence in interactive problem solving. This abstract does not claim to do that work; it is meant to provide some references to the literature and a starting-point for discussion.

Disclaimer

The 1997 AAAI Spring Symposium on qualitative decision theory revealed an emerging field with some promising ideas and potentially important applications. Unfortunately, there is still a large gap between the two. The most important trends that I can identify are still grappling in one way or another with the challenge of reworking a field that provides powerful theoretical arguments for representations of preferences that are not at all commonsensical, and that can be difficult to elicit.[1] At present, it is still unclear how to emerge from this foundational stage with an apparatus that will integrate with problem solving applications. The matter is complicated by the fact that there is no single dominant approach; different people have quite different ideas about how to proceed, and it may take a while for a few dominant paradigms to emerge.

Even if the foundations were much clearer, I think it would probably be far from obvious how to integrate decision-theoretic and planning approaches to decision making. Decision theory (even qualitative decision theory) can produce preferences over courses of action, but these preferences (which in general will take into account and resolve a large number of competing evaluative factors, as well as considerations having to do with risk) will not in general produce anything like a set of goals. But a planner needs goals in order to produce courses of action.

When I go to a (human) planning expert for advice, I take a number of preferences to the advice situation, some of them implicit and unexamined, and many of them competing. Suppose I am seeking financial planning. I am feeling restless and dissatisfied, and would rather retire early, though I will not be forced to retire on account of age. I am more timid about financial risk than I was when I was younger, and I'm unsure about the financial markets. I am increasingly uneasy about my health (this is one reason I want to retire early).

When I meet with the planning expert, somehow I focus on one part of this preference mix and produce a goal. I say "I want to retire at age 60." This conversation opener is clearly not a total picture of my state of mind. A good advisor may explore what happens when you plan on the supposed goal, but will also be quite willing to challenge the goal itself. Suppose that the expert runs some scenarios based on that goal and I am unhappy with them all. The expert asks me why I want to retire at 60. I find can't produce any very compelling reason. I chose the number 60 arbitrarily. My reasons for preserving my salary and employee benefits are much more compelling than my reasons for early retirement. The expert points out that a large proportion of early retirees are restless and dissatisfied with their retirement. At the end of the discussion, my preferences (or the weights I assign to preference factors) have changed. I decide to do nothing for at least five years, and to try harder to enjoy my job before rethinking the matter. I have been neither brainwashed nor compelled; the expert has helped me to reexamine and readjust my preferences.

It would be wonderful if we could replicate this sort of dynamic interaction between planning and preferences with an automated decision maker. But I suspect that this will require us to rethink not only the classical approach to decision theory, but the usual ap-

[1] I have in mind Savage's representation theorem in (Savage 1972), according to which the space of rational preferences over uncertain outcomes is isomorphic to a quantitative representation in terms of expected utility.

proaches to planning.

To return to the disclaimer. I wish that I could say more about how the following survey of work in qualitative decision theory is related to interactive problem solving. But I'm unable to do that, partly because I suspect that no one may have thought these issues through, and partly because I don't know enough about interactive and mixed-initiative decision making.

Background in QDT

Judging from the 1997 Spring Symposium on qualitative decision theory, there is still a certain amount of confusion and disagreement about research goals and basic theoretical ideas. Though (as in many other areas of AI) there are many quite different approaches and it is too early for us to be able to eliminate any of them with much confidence or to compare them formally, I think I can point to developing trends that will put the field on a much sounder footing.

One such trend is a separation of limited rationality issues from qualitative decision issues, and an appreciation of the potential importance of QDT for user modeling. At a strategic and foundational level, no one would doubt that qualitative decision making and limited rationality are closely related, but to try to address these larger issues at this point is only likely to impede progress. And (in ways that remind me of the development of nonmonotonic logic), I think many of us have abandoned the hope that qualitative methods will lead to much more efficient reasoning methods that in some sense are rational. The usefulness of relaxing the need for numerical representations lies not so much in any immediate efficiency payoff, but in eliminating the need to represent information that is often arbitrary and is difficult or impossible to obtain. The hope is that qualitative decision analyses will be more natural to think about, and that this will make things easier for system design and for user modeling. At the 1997 Spring Symposium, this point was made very convincingly by a number of people who had begun to consider user modeling issues.[2]

I am aware of the following five approaches to QDT: (1) Relaxations and specializations of quantitative expected utility, (2) Preference logics, (3) Approaches based on argumentation, (4) Approaches based on graphical representations, and (5) Approaches based on agent modeling. (These categories overlap to some extent.)

[2]See, for instance, (Bacchus & Grove 1997), (Boutilier *et al.* 1997), (Fox & Parsons 1997), and (Moorhouse & McKee 1997).

Relaxations and specializations

1. *Possibilistic logic* is an extension of Possibility Theory that takes utilities into account as well as probabilities. Possibility Theory, a theory inspired in part by fuzzy logic, is an approach to reasoning about uncertainty that has been under development for over ten years now; see (Dubois & Prade 1986). For a discussion of the extension to utilities, and of the problems involved in making this extension qualitative, see (Dubois, Prade, & Sabbadin 1997).

2. A plan p *dominates* a plan q when one can be sure that adopting p would yield better outcomes that adopting q.[3] Leonard Savage (Savage 1972) called the rule that a dominating action is to be preferred the *sure-thing principle.* As far as I know, the idea that purely qualitative dominance reasoning can be formalized, and used profitably in planning, is due to Michael Wellman (Wellman 1988).

(Thomason & Horty 1996) provides at least one attempt to provide a formal foundation for this idea. It develops a planning formalism similar to others that have been proposed for nondeterministic, multiagent planning. The models of this system resemble the decision trees of (Shafer 1997), but instead of assigning quantitative utilities to outcomes, the outcomes are merely compared by means of a partial order. In this approach, the arguments of the preference relation are *histories.* A history records the states that ensue as an agent executes the steps of a plan. In the simplest case, a history can be identified with the sequence of actions that makes up the plan. In the nondeterministic case, it will also record acts of nature and/or other agents.

By making independence assumptions concerning the actions of different agents, it is possible to define counterfactual selection functions in these models; we can perturb a history h in which a planning agent follows a plan p to obtain a set of *maximally close* histories in which the agent counterfactually follows another plan q. (Think of these as histories in which other agents act as they do in h.) These counterfactuals enable us to make sense of the phrase "would yield better outcomes" in the characterization of dominance. Formally, we can extend the preference relation over outcomes to a dominance relation over "histories." There is a soundness proof for this notion of dominance; it can be shown that any quantitative model can be soundly embedded in a quantitative decision-theoretic model.[4]

[3]This characterization of dominance is intentionally open-ended and informal.

[4]The quantitative theory is an alternative version of decision theory that—although it has received some attention in the foundational decision theoretic literature—is appar-

It has recently occurred to many members of the AI community that *multi-attribute utility*, as developed in works like (Keeney & Raiffa 1976) and (Fishburn 1977), can be an important component of a qualitative approach to decision theory. To my knowledge, the first citation of this work in the QDT literature is in (Doyle, Shoham, & Wellman 1991). Other papers that make use of multi-attribute utility are (Haddawy & Hanks 1993), (Draper, Hanks, & Weld 1994), (Linden, Hanks, & Lesh 1997), (Boutilier, Brafman, & Geib 1997), (Bacchus & Grove 1997), (Boutilier, Brafman, & Geib 1997), and (Shoham 1997).

Now, multi-attribute utility theory in itself is merely a way of representing and calculating the classical utilities of outcomes. The theory becomes relevant to qualitative (or at least nonclassical) approaches to decision analysis only when the decomposition of utilities into factors is combined with independence assumptions and with nonmonotonic estimations of utility. The papers mentioned above explore this idea; perhaps the most ambitious application of this idea is Yoav Shoham's work. At the moment, much of this work is unpublished, but there is some reason to hope that a network formalism for decision theory may emerge from it.

3. *Partial probability theory* relaxes classical probability by replacing a single probability measure by a family of such measures. A generalization of this idea to utility theory is discussed in (Voorbraak 1997).

Preference logics

1. Direct logics of preference have been studied in the philosophical literature; see, for instance, (Packard 1975), (Hansson 1990), and (Hansson 1996). Papers such as (Doyle & Wellman 1991), (Doyle, Shoham, & Wellman 1991), (Boutilier 1994), (Doyle & Wellman 1994), and (Doyle 1995) and (Bell & Huang 1997) present variations on this approach. Logics of this kind are compatible with most of the other approaches described here, and could serve as a useful way of organizing and comparing them.

Argument-Based approaches

1. There is a long and increasingly sophisticated tradition in modal logic that investigates conditional obligation, a notion that in many ways is similar to preference. This work in deontic logic began with Georg H. von Wright's work (see (von Wright 1963)), and in its more recent forms draws on ideas from nonmonotonic logic (see (Horty 1993) and (Horty 1994)). In the AI literature, there is a corresponding argument-based approach to decision-making. Some papers in this tradition include (Thomason 1993), (Fox & Parsons 1997), and (van der Torre & Tan 1997).

Graphical representations

1. It is very natural to try to extend the highly successful graphical representations that are used to model and implement probabilistic reasoning to deal with expected utility. Two closely related approaches of this sort are the *influence diagrams* that are have been used in medical domains (see (Nease, Jr. & Owens 1997) and (Owens, Schachter, & Nease, Jr. 1997)) and Shafer's decision trees (see (Shafer 1997)). Other uses of graphical representations in decision making include (Wellman 1988), (Bacchus & Grove 1997), and (Moorhouse & McKee 1997). I have already mentioned Shoham's work in this area, which strikes me as more general and foundationally important than existing approaches of which I'm aware.

Agent modeling

1. Yet another possible approach incorporates the agent models that are popular in the theory of distributed systems and other areas of AI, and seeks to extend them to include preferences. This approach is taken in (Brafman & Tennenholtz 1997).

Characterizing Sympathetic Planning Tasks

An adequate qualitative decision theory should provide a means of formalizing a fairly broad spectrum of decision problems (and this would certainly include cases of decision-making in the presence of uncertainty). It should yield solutions in many cases which, intuitively, can be solved without numerical calculation. Applications of multi-attribute utility theory to infer user preferences over fully determinate outcomes, as in (Linden, Hanks, & Lesh 1997) do not seem to require such an ambitious qualitative decision theory. In these cases, there is no element of uncertainty or risk in ranking choices, though—especially in the early stages—there may be uncertainty in the system's guesses about what these preferences are.

But there are cases in which a user seeks the help of an automated planning expert in order to get advice, not about alternative outcomes, but about alternative courses of action or plans. The user may hope to benefit from the computer's enhanced ability to search large plan spaces, and perhaps from its greater domain knowledge. But it goes without saying that the advice will not be appropriate unless it is based on an accurate model of the user's preferences. The task of

ently unfamiliar to most computer scientists. For information about this *causal* version of decision theory, see (Armendt 1988), (Gibbard & Harper 1978), (Eels 1982), and (Lewis 1981).

the system, then, is one of decision-theoretic planning, where beliefs (including probabilistic beliefs) are held constant, but where the utilities are adjusted to accommodate the user. This is what I want to call a *sympathetic planning task*.[5]

Sympathetic plan evaluation is needed not only in advisory tasks, but also in communicating plans. In (Young 1997), R. Michael Young points out that a full description of an automatic planner's output will have too much detail to be readily intelligible. In characterising what detail can be omitted, Young proposes that alternative plans that differ in ways that have little or no effect on the utility do not need to be distinguished by the plan description. For instance, if I am describing to a client how to get from the University of Pittsburgh to a hotel in Palo Alto, and I know that the client prefers the cheapest way of getting to the Pittsburgh airport, knows that the city airport bus is the cheapest way, and is familiar with the operation of the bus, I can simply say (1) rather than (2).

(1) Take the 3pm bus to the airport.

(2) Take the 3pm city transport bus to the airport. Go to the northeast corner of Fifth and Bigelow, arriving there by 2:55 pm. The bus may leave the corner before 3pm, but will never leave more than five minutes early. Have exact change for $.95 with you for the bus, or use a dollar bill if you prefer. The bus will make several stops en route to the airport. Stay on till the last stop, which is the airport.

Even the very verbose description in (2) depends on expectations about the user's knowledge and preferences. Some people, for instance, will prefer to be at the corner fifteen minutes early even if they are assured that the bus won't leave before 2:55. Description (2) strikes most people as unnatural (or even insulting and unnatural), and there is evidence that people understand terse plan descriptions like (1) much better than verbose ones. A terse plan description is appropriate if from it the hearer can reconstruct a course of action that will optimize the *hearer's*, not the speaker's preferences. So generating these terse descriptions will in general depend on a model of the user's preferences, as well as of the user's knowledge.

The components of sympathetic planning, then, are:

(i) The formation of a representation of another agent's utilities over an appropriate set of outcomes. This formation may well be achieved by an incremental improvement of partial representations.[6]

(ii) Generation of a set of alternative plans.

(iii) Evaluation of these alternatives using the represented utilities.

And if we want a sympathetic planner that can help its client to reassess goals, as in the retirement planning example I began with, we may need a fourth component.

(iv) Utility maintenance, or the ability to rationally revise preferences.[7]

I am not suggesting that these components should be separated in an implementation—in fact, they should certainly be interleaved. But they are conceptually distinct.

Some recent research trends in the decision-theoretic problem solving literature,[8] indicate growing interest in the problem of sympathetic planning.

The Challenge of Sympathetic Planning

Sympathetic planning is a natural and potentially useful extension of current planning systems, which (at least, in domains that involve uncertainty as well as substantial variation in the utilities of outcomes) offers the challenge of integrating the planning with a version of decision theory that is able to cope with risk. And even if a fully quantitative version of decision theory is used to evaluate the expected utility of plans, it seems very likely that a qualitative approach to utility would be needed in user modeling, and perhaps also in explaining the system's reasoning.

Therefore, a system of this kind will either have to integrate qualitative with quantitative expected utility, perhaps in the style of (Berleant & Kuipers 1997), or it will have to use qualitative decision-theoretic reasoning to evaluate plans, presumably reasoning that uses the same utility representations that are employed in user modeling.

Conclusion

It would be a joke to end this paper with a section titled "future work"—the whole paper is future work. An earlier version of this paper discussed the prospects for

[5]The word 'sympathy' usually connotes feeling what another person feels, if only in a diluted way. But to feel what another person feels, you need to form an impression of what they feel. So user modeling is a key part of sympathy.

[6]ha-etal:1997a,ha-haddawy:1997c make this point convincingly.

[7]A long-standing idea of Jon Doyles; see (Doyle 1980), for instance.

[8](Ha, Li, & Haddawy 1997), (Ha & Haddawy 1997), (Boutilier *et al.* 1997), and (Boutilier, Brafman, & Geib 1997).

using the dominance-based approach of (Thomason & Horty 1996) as a basis for sympathetic planning. But I think this idea is to preliminary to report on even in this forum.

Working out the ideas that are sketched above has helped me to remind me of the long-range goals I should keep in mind. I believe that the highest and best application of the ideas in qualitative decision theory that are outlined above is in the area of interactive planning, and in particular in sympathetic planning. it is very healthy for those of us who are struggling with the decision theoretic foundations to keep these goals in mind.

References

Armendt, B. 1988. Conditional preference and causal expected utility. In Harper, W. L., and Skyrms, B., eds., *Causation in Decision, Belief Change, and Statistics, Vol. 2.* Dordrecht: Kluwer Academic Publishers. 3–24.

B, C.; Brafman, R.; Geib, C.; and Poole, D. 1997. A constraint-based approach to preference elicitation and decision making. In Doyle, J., and Thomason, R. H., eds., *Working Papers of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, 19–28. Menlo Park, California: American Association for Artificial Intelligence.

Bacchus, F., and Grove, A. 1997. Independence in qualitative decision theory. In Doyle, J., and Thomason, R. H., eds., *Working Papers of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, 1–8. Menlo Park, California: American Association for Artificial Intelligence.

Bell, J., and Huang, Z. 1997. Dynamic goal hierarchies. In Doyle, J., and Thomason, R. H., eds., *Working Papers of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, 9–17. Menlo Park, California: American Association for Artificial Intelligence.

Berleant, D., and Kuipers, B. J. 1997. Qualitative and quantitative simulation: Bridging the gap. *Artificial Intelligence* 95(2):215–255.

Boutilier, C.; Brafman, R.; and Geib, C. 1997. Prioritized goal decomposition of Markov decision processes: Towards a synthesis of classical and decision theoretic planning. In Pollack, M., ed., *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence.* San Francisco: Morgan Kaufmann.

Boutilier, C. 1994. Defeasible preferences and goal derivation. Unpublished manuscript; Department of Computer Science, University of British Columbia, Vancouver, British Columbia, Canada V6T 1Z4.

Brafman, R., and Tennenholtz, M. 1997. On the axiomatization of qualitative decision criteria. In Doyle, J., and Thomason, R. H., eds., *Working Papers of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, 29–34. Menlo Park, California: American Association for Artificial Intelligence.

Doyle, J., and Wellman, M. 1991. Preferential semantics for goals. In Dean, T., and McKeown, K., eds., *Proceedings of the Ninth National Conference on Artificial Intelligence*, 698–703. Menlo Park, California: American Association for Artificial Intelligence.

Doyle, J., and Wellman, M. 1994. Representing preferences as *ceteris paribus* comparatives. In *Working Notes of the AAAI Spring Symposium on Decision-Theoretic Planning*, 69–75. Menlo Park, California: American Association for Artificial Intelligence.

Doyle, J.; Shoham, Y.; and Wellman, M. P. 1991. A logic of relative desire (preliminary report). In Ras, Z., ed., *Proceedings of the Sixth International Symposium on Methodologies for Intelligent Systems*, Lecture Notes in Computer Science. Berlin: Springer-Verlag.

Doyle, J. 1980. A model for deliberation, action, and introspection. Technical Report AI TR 581, Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Doyle, J. 1995. Contextual equivalence and *ceteris paribus* comparatives. Unpublished manuscript, Laboratory for Computer Science, Massachusetts Institute of Technology.

Draper, D.; Hanks, S.; and Weld, D. 1994. Probabilistic planning with information gathering and contingent execution. In Hammond, K., ed., *Proceedings of the Second International Conference on AI Planning Systems*, 31–36. Menlo Park, California: American Association for Artificial Intelligence.

Dubois, D., and Prade, H. 1986. *Possibility Theory: An Approach to Computerized Processing of Uncertainty.* New York: Plenum Press.

Dubois, D.; Prade, H.; and Sabbadin, R. 1997. A possibilistic logic machinery for qualitative decision. In Doyle, J., and Thomason, R. H., eds., *Working Papers of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, 47–54. Menlo Park, California: American Association for Artificial Intelligence.

Eels, E. 1982. *Rational Decision and Causality.* Cambridge, England: Cambridge University Press.

Fishburn, P. C. 1977. Multiattribute utilities in expected utility theory. In Bell, D. E.; Keeney, R. L.; and Raiffa, H., eds., *Conflicting Objectives in Decisions.* New York: John Wiley and Sons. 172–196.

Fox, J., and Parsons, S. 1997. On using arguments for reasoning about actions and values. In Doyle, J., and Thomason, R. H., eds., *Working Papers of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, 55–63. Menlo Park, California: American Association for Artificial Intelligence.

Gibbard, A. F., and Harper, W. L. 1978. Counterfactuals and two kinds of expected utility. In Hooker, C.; Leach, J.; and McClennen, E., eds., *Foundations and Applications of Decision Theory.* Dordrecht: D. Reidel Publishing Co. 125–162.

Ha, V., and Haddawy, P. 1997. Problem-focused incremental elicitation of multi-attribute utility models. In *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence (UAI-97)*, 215–222.

Ha, V.; Li, T.; and Haddawy, P. 1997. Case-based preference elicitation (preliminary report). Unpublished manuscript, Department of Electrical Engineering and Computer Science, University of Wisconsin-Milwaukee.

Haddawy, P., and Hanks, S. 1993. Utility models for goal-directed decision-theoretic planners. Technical Report 93-06-04, Department of EE&CS, University of Wisconsin-Milwaukee, Milwaukee, Wisconsin.

Hansson, B. 1990. Preference-based deontic logic. *Journal of Philosophical Logic* 19:75–93.

Hansson, S. O. 1996. What is *ceteris paribus* preference? *Journal of Philosophical Logic* 25(3):307–332.

Horty, J. F. 1993. Deontic logic as founded on nonmonotonic logic. *Annals of Mathematics and Artificial Intelligence* 9:69–91.

Horty, J. F. 1994. Moral dilemmas and nonmonotonic logic. *Journal of Philosophical Logic* 23(1):35–65.

Keeney, R. H., and Raiffa, H. 1976. *Decisions With Multiple Objectives: Preferences and Value Tradeoffs.* New York: John Wiley and Sons, Inc.

Lewis, D. K. 1981. Causal decision theory. *Australasian Journal of Philosophy* 59:5–30.

Linden, G.; Hanks, S.; and Lesh, N. 1997. Interactive assessment of user preference models: The automated travel assistant. In *Proceedings, User Modelling '97.* Available at www.cs.washington.edu/homes/hanks/Papers/.

Moorhouse, J., and McKee, G. 1997. Modelling evaluative judgements. In Doyle, J., and Thomason, R. H., eds., *Working Papers of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, 71–79. Menlo Park, California: American Association for Artificial Intelligence.

Nease, Jr., R. F., and Owens, D. K. 1997. Use of influence diagrams to structure medical decisions. *Medical Decision Making* 17:263–275.

Owens, D. K.; Schachter, R. D.; and Nease, Jr., R. F. 1997. Representation and analysis of medical decision processes with influence diagrams. *Medical Decision Making* 17:241–262.

Packard, D. J. 1975. A preference logic minimally complete for expected utility maximization. *Journal of Philosophical Logic* 4(2):223–235.

Savage, L. 1972. *The Foundations of Statistics.* New York: Dover, 2 edition.

Shafer, G. 1997. *The Art of Causal Conjecture.* Cambridge, Massachusetts: The MIT Press.

Shoham, Y. 1997. A mechanism for reasoning about utilities (and probabilities): Preliminary report. In Doyle, J., and Thomason, R. H., eds., *Working Papers of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, 85–93. Menlo Park, California: American Association for Artificial Intelligence.

Thomason, R. H., and Horty, J. F. 1996. Nondeterministic action and dominance: Foundations for planning and qualitative decision. In Shoham, Y., ed., *Theoretical Aspects of Rationality and Knowledge: Proceedings of the Sixth Conference (TARK 1996).* San Francisco: Morgan Kaufmann. 229–250. The statement and proof of the soundness theorem in Section 7 of this version are flawed. See www.pitt.edu/~thomason/dominance.html.

Thomason, R. H. 1993. Towards a logical theory of practical reasoning. In *Working Notes of the AAAI Spring Symposium on Reasoning about Mental States.* Menlo Park, CA: American Association for Artificial Intelligence.

van der Torre, L. W., and Tan, Y.-H. 1997. Prohairetic deontic logic and qualitative decision theory. In Doyle, J., and Thomason, R. H., eds., *Working Papers of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, 103–111. Menlo Park, California: American Association for Artificial Intelligence.

von Wright, G. H. 1963. *Norm and Action.* London: Routledge and Keegan Paul.

Voorbraak, F. 1997. Decision analysis using partial probability theory. In Doyle, J., and Thomason, R. H., eds., *Working Papers of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, 113–119. Menlo Park, California: American Association for Artificial Intelligence.

Wellman, M. P. 1988. Formulation of tradeoffs in planning under uncertainty. Technical Report TR–427, Massachusetts Institute of Technology Laboratory for Computer Science, 545 Technology Square, Cambridge Massachusetts, 02139.

Young, R. M. 1997. *Generating Concise Descriptions of Complex Activities.* Ph.d. dissertation, Intelligent Systems Program, University of Pittsburgh, Pittsburgh, Pennsylvania.