Critique of KRSL 2.0.2

Jon Doyle
MIT Lab for Computer Science
doyle@lcs.mit.edu
October 11, 1994

Abstract

The current KRSL plan specification language treats information traditionally contained in plans reasonably well, but does not provide good means for representing economic information about how agents consume, produce, and choose among resources, and provides essentially no means for representing information about the rationales (reasons, decisions, and explanations) underlying plans. In consequence, planning agents cannot use KRSL to cooperatively make decisions and evaluate and repair plans, as they cannot use KRSL to communicate useful information about their needs (in order to evaluate plans) and reasons (in order to modify plans).

The KRSL plan-representation language is intended to provide a common language to ARPI groups for specifying information about plans and communicating such information among modules. The clear benefit of such a common language is that it simplifies the task of integrating modules by getting them all to speak the same language, rather than having to design a large numbers of point-to-point translators of idiosyncratic representations. Unfortunately, the current version KRSL facilitates communication of only a portion of the information needed by mixed-initiative planning systems. It focuses on describing the overt structures of objects and actions in a way completely appropriate to traditional AI planning systems. But it provides flawed means for specifying economic properties of agents and resources, and provides no standard means for communicating plan rationales.

In the original ARPI vision of human-computer cooperative planning, human planners would work with automated systems to present commanders with a short list of primary, qualitatively different options, the idea being to give the commanders a good idea of the range of possibilities seen as reasonable by the planners. To allow the commander(s) to choose wisely, these options must be accompanied by at least a sketch of their underlying rationales if the options are not absolutely transparent (as is rarely the case). These rationales indicate the primary planning assumptions and decisions making the options plausible; the underlying assumptions about conditions and responses, the underlying assumptions about the relative importance of different strategic and tactical aims and principles, and the principal choices among equally good alternatives (equivalently, the dimensions of simple variability). One plan might presume that some position will hold without fail, or might chance greater risks than another, or might sacrifice one goal to accomplish another. Commanders need to know these things about the options to decide among them. For similar reasons, replanning systems designed to repair or revise large-scale plans in light of changed information also need the rationales of plans if they are to do anything other than replanning from scratch.

The current KRSL specification provides no standard means for communicating useful plan rationales. One can state that a plan achieves a specific goal in KRSL; but one cannot say much about underlying assumptions about conditions, or the alternatives considered in selecting this plan over others that achieve the goal, or about the preferences that lead to this selection, or about the importance of this goal to others. Accordingly, KRSL should be improved to permit convenient specification of plan rationales and the information contained in them. Past work on reason maintenance [2, 5] and reasoned deliberation [3] provides suggestions for some elements of rationales, namely argument fragments and descriptions of decisions.

Improving KRSL's language of plan rationales in these ways will directly benefit the designers of intermediate-level planning system components (as in the ACP task), who will use this portion of KRSL to communicate rationale information crucial in revising plans. Plan rationales indicate which parts of a plan are separate from each other, or more generally, the dependence of parts on others, and thus provide the basic information needed to minimize the portion of a plan needing repair. But rationales also include the preferential judgments or decisions used in choosing one alternative part over others, and thus provide both immediate alternatives to consider in revisions, as well as providing strategic information about what revisions are more important than others. The current KRSL, in focusing on communicating only plans but not rationales, provides only plans which modules may take or leave as wholes, with no indication of component interdependence or importance.

One key element of plan rationales present in KRSL but treated inadequately is preferential and utility information. The language includes a construct for specifying preferences among things, but without any clear meaning given to the things so related. The descriptive text seems to indicate that preferences relate propositions (sets of possible worlds) or descriptions of such, and mean that every world satisfying one proposition or description is preferred (perhaps weakly) to every world satisfying the other. This semantics is pretty useless in practice, for it forbids distinguishing tradeoffs preferentially—that is, all outcomes trading off mutually unsatisfiable conditions must be indifferent. The language includes means for stating utility functions, but provides no structure for stating the standard constructs of multiattribute utility theory [9]. Accordingly, KRSL should be corrected to provide a library of standard constructs from multiattribute utility theory making convenient specification and composition of subutility functions [13], and to draw on recent work on stating propositional preferences correctly [1, 4, 6, 10, 11, 12]. It should also incorporate standard constructs relating planning goals to preference and utility measures, including both the theoretical constructs just mentioned, and libraries like those of Haddawy and Hanks [7, 8] identifying specific utility functions approximating common types of planning and problem-solving goals.

Preferences constitute only one form of economic information treated inadequately by KRSL. The current language provides rudimentary means for specifying production and consumption of resources, but little is expressible beyond the simple fact of production or consumption. Standard economic theory provides additional concepts for systematizing specifications of economic behaviors through its general and special notions of production functions, which indicate how consumption depends on production and vice versa. In addition, standard economic theory provides a taxonomy of different types of consumers, producers, resources, and auction methods useful in describing market and choice mechanism, as well as descriptors of the structure of production functions, such as "complementarities" (necessary coordinations among resources, e.g., guns aren't much use without bullets and vice versa).

Put together, these improvements to KRSL's economic language should prove very useful to communicating plans and their rationales. Planning experts working to codify planning problems will find the improved and expanded language significantly more useful in accurately characterizing planning situations, resources, and tasks, as the language will finally be able to express in convenient and standard ways both the ordinary production/consumption relationships among actors and resources and the strategic importance among different resources and activities within a plan. For example, the current KRSL provides only indirect ways of saying that tanks need ammunition and vice versa, or that transporting more fighter pilots is more important than transporting more press officers. The improved language should make accurate expression of such information possible, and will facilitate the sharing and communication of this information among a wide range of systems, both human and machine.

Some other problems include the following.

- There is no way to refer to the context in which statements are valid except temporal contexts.
- Resource usage records just seem to indicate current usages. Can they also indicate schedules of usages under different conditions?
- Can quantities be schedules? (Mappings?) Mappings from contexts to quantities?
- There are no assumptions (nonmonotonic) apart from defaults in objects. Do concepts have no default values?
- There are no handles on representations themselves.

References

[1] C. Boutilier. A modal characterization of defeasible deontic conditionals and conditional goals. In *AAAI Spring Symposium on Reasoning about Mental States*, pages 30–39, 1993.

[2] J. Doyle. A truth maintenance system. *Artificial Intelligence*, 12(2):231–272, 1979.

[3] J. Doyle. A model for deliberation, action, and introspection. AI-TR 581, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, 545 Technology Square, Cambridge, MA, 02139, 1980.

[4] J. Doyle, Y. Shoham, and M. P. Wellman. A logic of relative desire (preliminary report). In Z. W. Ras and M. Zemankova, editors, *Methodologies for Intelligent Systems, 6*, volume 542 of *Lecture Notes in Artificial Intelligence*, pages 16–31, Berlin, Oct. 1991. Springer-Verlag.

[5] J. Doyle and M. P. Wellman. Rational distributed reason maintenance for planning and replanning of large-scale activities. In K. Sycara, editor, *Proceedings of the DARPA Workshop on Planning and Scheduling*, pages 28–36, San Mateo, CA, Nov. 1990. Morgan Kaufmann.

[6] J. Doyle and M. P. Wellman. Representing preferences as *ceteris paribus* comparatives. In S. Hanks, S. Russell, and M. P. Wellman, editors, *Proceedings of the AAAI Spring Symposium on Decision-Theoretic Planning*, 1994.

[7] P. Haddawy and S. Hanks. Issues in decision-theoretic planning: Symbolic goals and numeric utilities. In *Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control*, pages 48–58, 1990.

[8] P. Haddawy and S. Hanks. Representations for decision-theoretic planning: Utility functions for deadline goals. In B. Nebel, C. Rich, and W. Swartout, editors, *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning*, pages 71–82, San Mateo, CA, 1992. Morgan Kaufmann.

[9] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Trade-offs.* John Wiley and Sons, New York, 1976.

[10] S. Mantha. Towards a logical foundation for qualitative decision theory. In S. Hanks, S. Russell, and M. P. Wellman, editors, *Working Notes of the Symposium on Decision-Theoretic Planning*, pages 169–174. American Association for Artificial Intelligence, Mar. 1994.

[11] S.-W. Tan and J. Pearl. Specification and evaluation of preferences for planning under uncertainty. In J. Doyle, E. Sandewall, and P. Torasso, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourth International Conference (KR'94)*, San Francisco, CA, May 1994. Morgan Kaufmann.

[12] M. P. Wellman and J. Doyle. Preferential semantics for goals. In *Proceedings of the National Conference on Artificial Intelligence*, pages 698–703, 1991.

[13] M. P. Wellman and J. Doyle. Modular utility representation for decision-theoretic planning. In *Proceedings of the First International Conference on AI Planning Systems*, 1992.