Robotic Assembly and Task Planning

Avi Kak

If classical planners are ever to automatically plan the actions of the smart machines, particularly robots for the automatic assembly of industrial objects, then they will have to know much more about geometry and topology as well as sensing. Consider that the simple act of changing an object's grasp-the change might be necessitated by the nature of some assembly goal-involves the interaction of the geometries of the grasping device and the object if the change is to occur without a collision between the device and the object. Of course, one could ask, Could geometric considerations be divorced from the highly developed symbolic-level planning? That is, could we first synthesize a symbolic plan and then plug in the geometry for the execution of the actions? Experience has shown the answer to be, unfortunately, a big no. The actions chosen to fulfill an assembly goal depend as much on the symbolic nature of the goal-for example, by the form (PUTON A B)—as they do on the geometric constraints involved in putting object A on top of object B-for example, we might want certain features on the exteriors of A and B to line up as a result of the assembly action.

Then, there is the ever-present issue of how to integrate sensing with planning. This integration must be achieved if machines are to be able to use sensing to reduce the uncertainties in their environments. This issue is complex for a variety of reasons: An action can either increase or decrease the uncertainties associated with the environment and do so in a differential manner; that is, an action can cause a decrease in the uncertainty associated with, say, the x coordinate of the position of an object while causing an increase in the

uncertainty associated with the y coordinate. A sensor has to be chosen in light of this interplay between actions and uncertainties. The planner must be able to figure out what actions are executable given the current uncertainties and what are executable if appropriate sensors are invoked to reduce these uncertainties.¹ In an attempt to answer these issues, the planner Spar was developed in the Robot Vision Lab at Purdue University. Now implemented on a Puma 762 robot, Spar is discussed in "Spar: A Planner That Satisfies Operational and Geometric Goals in Uncertain Environments" by Seth A. Hutchinson and Avinash C. Kak.

Input to planners such as Spar consists of assembly goals; these goals specify what two parts are to come together, the initial and the final positions of these parts, and the uncertainties associated with their initial placement. Who should specify these assembly goals? Should they be supplied by a human expert? Should they be discovered by an automated reasoning agent starting from, say, a computer-aided design description of the assembled product? In the well-used example of flashlight assembly, some intelligent agent has to discover the fact that the two end pieces cannot be installed on the tube of the flashlight before the battery is inserted. In "Assembly Sequence Planning," Arthur C. Sanderson, Hui Zhang, and Luiz Homem de Mello show how geometric reasoning can be carried out in the Pleiades system to discover optimum sequences of assembly operations with respect to such criteria as geometric feasibility, attachment feasibility, and tool availability.

Clearly, the methods for determining geometric feasibility are important not only at assembly design time but also when optimum assembly sequences are computed. How objects are represented, particularly with regard to their symmetries, has great bearing on the methods that can be invoked for testing the geometric feasibility of an assembly. As Robin Popplestone, Yanxi Liu, and Rich Weiss discuss in "A Group Theoretic Approach to Assembly Planning," group theory provides us with a powcrful and compact formalism for representing the symmetries of objects. In addition, it gives us a set of tools for inferring the overall symmetries of an object assembled from parts whose symmetries are already known. Furthermore, if an assembly consists of multiple features coming together from different objects, then group theory can also be used to test the geometric feasibility of the overall assembly by intersecting the constraints corresponding to each pair of mating features.

This special issue on robotic assembly and task planning consists primarily of these three articles. My wish was to include one more, "AI Planning Systems: Problems and Solutions" by Austin Tate, James Hendler, and Mark Drummond; however, a combination of circumstances precluded its publication in this issue. A general survey of the literature on classical planning, this article will appear in the next issue of *AI Magazine*.

Notes

1. Of course, not to be forgotten is the use of sensors to endow a robot with reflexive behavior. In the context of robotic assembly, such behavior is important for tasks such as mating parts under tight tolerances using compliant motions generated by force-torque feedback.

Letters

Editor:

"Fronti Nulla Fides": No reliance can be placed on appearance.

In a letter to the editor (*AI Magazine*, Winter 1989), Benjamin Kuipers criticizes various points made in an earlier paper of ours (Akman and ten Hagen 1989).

First, a side (nonetheless important) remark: Although Kuipers asserts that he distributes QSIM to interested researchers, our experience has been otherwise. Akman has tried twice to obtain QSIM, without success. Although Kuipers promised to deliver a copy—QSIM was under revision at the time of Akman's request (this being as early as winter 1988) the program was never sent. So much for the availability of QSIM....

Kuipers' letter is full of sweeping generalizations that are so much against the nature of scientific enterprise. We should also add that we are disappointed to see Kuipers employing universal truths and unarguable facts such as ". . . if you build the wrong model, the predictions derived from that model are likely to be wrong" or "... guarantees of mathematical validity [are] necessary for any science" as his main cheval de bataille. In the following we'll point out, one by one, the weaknesses of QSIM. Our task will be easy since we shall merely reproduce, almost verbatim, Kuipers' own sentences (Kuipers 1986) and, additionally, Janowski's (1987) views. (The latter reference gives an excellent review of QSIM's disadvantages.) Then, we'll let the reader judge.

... the QSIM algorithm, and local qualitative simulation algorithms in general, cannot be guaranteed against producing spurious behaviors; behaviors which are not actual behaviors for *any* physical system satisfying the constraint equations [Kuipers' italics] (Kuipers 1986, pp. 317–318).

... if a valid description of the mechanism can produce invalid predictions (false positives), its use-

fulness is limited [our italics] (Kuipers 1986, p. 318).

If we explicitly add the constraints representing conservation of energy to the oscillating spring constraint equations, the single correct behavior is found. However, although the additional constraints are derivable from the original equations, *it is not at all clear how to do such a derivation automatically for an arbitrary mechanism* [our italics] (Kuipers 1986, p. 321).

If qualitative simulation yields several possible behaviors, further analysis is required before concluding that they represent possible futures [our italics] (Kuipers 1986, p. 321).

QSIM is just such a method of fitting together the jigsaw of the histories of individual parameters, by constraint propagation, in order to derive the possible behaviors of a system [our italics] (Janowski 1987, p. 67).

For example, it is possible to simulate the behavior of an oscillating spring. But without information about dissipation of energy, QSIM is unable to tell whether each successive oscillation is greater, smaller, or the same as the previous one. Accordingly, it generates all three possible behaviors. Obviously, only the decreasing oscillations represent a real-world behavior, and the others are spurious. This illustrates an important point about QSIM: if it is given a correct description of the real-world, then all real-world behaviors of the system will be simulated, but not all the simulated behaviors will necessarily be possible in the system [our italics] (Janowski 1987, p. 69).

It is not the role of QSIM to create the initial model: QSIM is a method of solving the problems once they have been formulated. By contrast, [Ken] Forbus, and [Johan] de Kleer and [John Seely] Brown try to define a "physics," of which generating behavioral descriptions is one part but in which they also attempt to formalize the translation from the real world to the model [our italics] (Janowski 1987, p. 70).

Thus, there is a distinction between, on the one hand, a "physics" based on general principles which has formalisms for modeling the real world in a qualitative manner and, on the other hand, a tool for "handcrafting" such models more or less individually. It is in the second camp that QSIM sits [our italics] (Janowski 1987, p. 70).

We agree that QSIM is a mathematical language for expressing models, rather than a physical methodology for building correct models. However, this doesn't imply, as Kuipers suggests, that a modeling language must be able to express both good and bad models, both true and false models. This is equivalent to saying that all the responsibility for guaranteeing good models (or true models) is delegated to some outside agent (a person?). We then ask, Where is the intelligence of QSIM if one needs such a drastic amount of intervention?

Kuipers' letter also tries to establish a point which sounds plausible to us and probably to many people working in the area; viz., QSIM is a step toward providing the degree of expressive and inferential power necessary for qualitative physics. Yes, but it is one step! To paint a rosy picture around an "all-powerful QSIM" is an attempt at creating a myth and is dangerous. D. McDermott said, "A common idiocy in AI research is to suppose that having identified the shortcomings of Version I of a program is equivalent to having written Version II" (McDermott 1981). The shortcomings of QSIM are not of the kind that can be corrected in a second version. (QSIM's precursor was ENV [Kuipers 1984].)

Finally, as for the simplistic remarks of Kuipers about causality, we suggest

that he take a look at the modern accounts of causation as presented, say, in Shoham (1987).

Probably, it is apt to conclude this discussion on a lighter note. To quote McDermott again, "To say anything good about anyone is beyond the scope of this letter."

Varol Akman Bilkent University, Ankara Paul ten Hagen Center for Mathematics and Computer Science, Amsterdam

References

Akman, V., and ten Hagen, P. 1989. The Power of Physical Representations. AI Magazine 10(3): 49-65.

Kuipers, B. 1986. Qualitative Simulation. Artificial Intelligence 29:289-338.

Kuipers, B. 1984. Commonsense Reasoning about Causality: Deriving Behavior from Structure. Artificial Intelligence 24:169-203.

Shoham, Y. 1987. Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence. Cambridge, Mass.: MIT Press.

Janowski, R. 1987. An Introduction to QSIM and Qualitative Simulation. Artificial Intelligence in *Engineering* 2(2):65–71.

McDermott, D. 1981. Artificial Intelligence Meets Natural Stupidity. In Mind Design, ed. J. Haugeland, 143-160. Montgomery, Vt.: Bradford.

This publication is available in microform from University Microfilms International.



Call toll-free 800-521-3044. In Michigan, Alaska and Hawaii call collect 313-761-4700 Or mail inquiry to: University Microfilms International, 300 North Zeeb Road, Ann Arbor, MI 48106

For free information, circle no. 116

Research and Development Opportunities in Knowledge-Based Systems Technology

The Price Waterhouse worldwide financial services business provides a wide range of opportunities for using knowledge-based systems technology. In order to be realized, many of those opportunities require significant advances in the underlying technologies. The Price Waterhouse Technoloav Centre provides an environment in which high quality research and development professionals can pursue those advances and explore the strengths and weaknesses of the resulting technology by using it to build prototype systems that address central problems in the firm's business. The Centre is located in Menlo Park, California, near Stanford University. The Knowledge Based Systems group, directed by Richard Fikes, is currently working on representation and reasoning techniques that support modeling of businesses and tax regulations, and the use of those models for diagnosis, planning, and risk assessment.

If you have an established track record in research and/or the application of advanced AI techniques, and you wish to join an environment that encourages technical creativity and professional arowth, we invite you to consider joining us.

Please contact us at:

Price Waterhouse Technology Centre

Human Resources Department 68 Willow Road Menlo Park, CA 94025

podgorsk@pw.com Fax: (415) 321-5543

An equal opportunity affirmative action employer

Price Waterhouse



Expect more from us