

From: Proc. of the 6th European Meeting on Cybernetics and Systems Research,
Univ. of Vienna, Vienna, Austria, Apr. 13-16, 1982.

A PLANNING CONTROL STRATEGY
THAT ALLOWS FOR THE COST OF PLANNING

Daniel T. Johnson and Lenhart K. Schubert*

Department of Computing Science
University of Alberta, Edmonton, Alberta, Canada T6G 2H1

This paper addresses the problem of optimizing the cost of a plan, including both the planning and execution costs. The fundamental difficulty in formulating an optimal strategy is the fact that the relationship between the cost of a goal and the costs and probabilities of its subgoals is itself a function of the strategy used to search these subgoals. The proposed escape from the difficulty relies on approximate prediction of how the strategy will continue searching the subgoals of a chosen goal once it has made its initial choice among these subgoals. The strategy has been proved theoretically optimal for a certain version of the multi-box monkey-and-bananas problem, and has been successfully hand-simulated in a STRIPS-type robot's world of fairly realistic magnitude.

1. Introduction

The need to allow for the cost of planning in controlling a planning process is apparent even for simple problems. Consider the monkey-and-bananas problem, generalized so that many alternative boxes scattered about the monkey's cage could be used by the monkey to reach the bananas. By pondering the cost of using each box, the monkey can formulate a plan whose expected execution cost is minimal; this is the type of plan optimization considered by Feldman & Sproull (1974). However, thinking is not free; if the number of available boxes is large, the monkey will spend more time 'optimizing' his choice of a box than he would take to reach the bananas using the very first box he sees. Feldman & Sproull recommend abandoning efforts to improve a plan and switching to execution mode when the cost of the next planning step exceeds the difference between the maximum and expected utility of the plan. However, this criterion is unusable in a hierarchical planning situation, where 'improving a plan' is not clearly distinct from merely proceeding with its elaboration. Besides, the effect of the criterion depends entirely upon the choice of planning 'quantum': if planning effort

is allocated in sufficiently small quanta, then additional planning will always win over execution as long as there is any potential for improvement. This is the same type of criticism that is often levelled at optimizing strategies such as A* (see Nilsson, 1980). For example, a robot using A* to collect a few hundred nuts and bolts spilled on the floor would be immobilized by the explosively complex problem of optimizing the sequence of moves for accomplishing the task.

The strategy we will propose resembles that of MULTIPLE (Slagle & Bursky 1968) in that it relies on characterizations of subgoals which are supplied by heuristic functions at the tip nodes of the current planning tree and are "backed up" to the top level whenever new subgoals are sprouted. The characterizations of subgoals involve cost estimates (i.e., estimates of the cost of completing and executing a plan to achieve the goal), success probability estimates, and approximate distributions describing the changes in cost and probability estimates which may result from an initial investment of effort in the subgoals. We argue that such multiparameter characterizations are unavoidable in general.

The strategy makes intuitively correct decisions in cases where pure

* Current address of first author: 4713 N. Cashel Circle, Houston, Texas 77069. Please address reprint requests to second author.

utility-based or probability-based methods go astray. It is in some cases provably optimal; specifically, we have established its theoretical optimality in a version of the multi-box monkey-and-bananas problem (recast somewhat more realistically as the problem of purchasing a low-cost item such as a snow shovel or a used tricycle through the want ads in a city newspaper). We have also carried out two successful hand-simulations of the strategy in a STRIPS-type robot's world. (For full details of the present work see Johnson, 1977).

2. Allowing for the cost of planning

The cost of planning has two components: the cost of expanding a non-primitive goal into a sequence of one or more simpler subgoals (i.e., transforming the goal or forming a plan for achieving it), and the cost of selecting a goal for expansion.

We propose to take formal account of the cost of goal expansion by lumping it together with the cost of execution. In other words, we think of the cost of 'executing' a higher-level goal or plan as including the cost of fully expanding it into a sequence of primitive steps as well as the cost of carrying out those primitive steps. Of course, goal costs will then depend on the planning control strategy; however, for a system operating with a fixed control strategy, there should be no more difficulty in formulating (and adaptively modifying) heuristic estimators for these comprehensive costs than in formulating them for execution costs alone. The real difficulty lies in determining how the costs of alternative subgoals of a given goal combine, since this combined cost is strategy-dependent. This problem has therefore been the main focus of our investigation.

Can the goal selection cost be added to the execution cost as well? Apparently not, since that cost is context-dependent. For example, if some goal G is the only available choice at some point in a planning process, the cost of selecting it will be lower than if there are many alternative choices requiring comparative evaluation. Attempts to allow for goal selection costs in the goal selection strategy are thwarted by an even more insidious difficulty: anticipating those costs will itself be more or less costly, so that an optimal strategy should also anticipate the cost of estimating the cost of goal selection; and so on ad

infinitum. This is one point at which we encounter the classic 'cost of analysis' problem of decision theory.

However, the difficulties concerning goal selection cost become unimportant if we can assume that the goal selection cost constitutes a relatively small fraction of the total planning cost, i.e., that goal selection is cheap compared to goal expansion and execution. If we proceed on that assumption and arrive at a control strategy which is indeed computationally inexpensive, then the assumption is vindicated.

We now turn to the required characterization of subgoals. In the first place, we require a total cost estimate at each subgoal, i.e., an estimate of the cost of fully elaborating a plan (to the level of immediately executable actions) and executing it. The need for such cost estimates is easily appreciated for problems such as the multi-box monkey-and-bananas problem, where the success probability may be unity for all subgoals, so that a purely probability-based method like MULTIPLE is left with no basis for preferring one over another. On the other hand probabilities seem to be indispensable as well, despite attempts such as that of Feldman & Sproull (1974) to parameterize planning problems in terms of utility alone. As will be seen below, the correct measure of 'worthiness' of alternative goals in certain simple cases is their ratio of success probability to cost; this is obviously not a monotonic function of utility (except at constant probability) and so in general will produce a different goal ranking than utility. Nor is that measure of 'worthiness' sufficient by itself, since a different measure is needed for conjoined goals, namely the ratio of failure probability to cost.

But even cost and probability estimates are inadequate for rational planning in general. Consider another variant of the monkey-and-bananas problem, in which the monkey is next to a box which may be either very easy (cheap) or very hard (expensive) to push. He must choose between this box and a second, more distant box whose cost has a known, intermediate value. Even though the expected cost of the first box may be higher than that of the second (perhaps even infinite, if the box might be bolted to the floor), it may well be to the monkey's advantage to try the first box initially, on the

chance that it may be very easy to push. Thus, a higher-expected-cost plan may be preferable to a lower-expected-cost plan if it offers some chance of a very cheap solution (and can be abandoned prematurely if it proves unprofitable). Similarly, a lower-probability plan may be preferable to a higher-probability plan, even if its expected cost is no lower, in a case where a small initial investment may greatly improve the chance of its success. (This intuition presumably motivated the use of a measure of 'merit', or absolute rate of change of probability, in MULTIPLE.) For these reasons, and because of the specific needs arising from the attempt to formulate a strategy which anticipates its own behaviour, we have chosen a subgoal characterization which supplies not only the current cost and probability estimates, but also distributions over the anticipated changes in these estimates after an initial investment of effort. These characterizations are assumed to be supplied by heuristic functions at the tip nodes of the planning tree, and are backed up to the non-terminal nodes to allow selection of a globally optimal subgoal.

3. The goal selection strategy

Suppose that a goal G of a plan can be pursued in either of two ways, namely by pursuing goal G1 or goal G2. Suppose further that G1 and G2 are tip nodes of the current planning tree and that we have heuristic estimates (p_1, t_1) , (p_2, t_2) of their probability and cost. The cost estimate for a tip node reflects the expected cost of pursuing the corresponding goal exclusively until it succeeds or fails, with failure becoming evident after a finite investment of effort. We now wish to characterize the probability and cost of G if G is pursued exclusively but the choice between G1 and G2 is subject to reconsideration. In the same way we then want to back up characterizations from G and its "siblings" to G's parent, assuming exclusive commitment to G's parent but allowing for reconsideration of the choice among G and its siblings; and so on. Each backing-up step is also required to compute and record the choice of subgoal to which the next unit of effort should be devoted (much as in MULTIPLE). Eventually we should arrive at a characterization of the top goal which presupposes an exclusive commitment to this goal but allows fully for reconsideration of lower-level choices. The ability to reconsider choices is of course crucial to a goal

selection strategy; it must be possible to abandon a selected goal in favour of an alternative if an initial attempt to expand it indicates that it is much costlier or less likely to succeed than previously anticipated.

Optimizing the initial choice between G1 and G2 in a way which allows for subsequent reconsideration of the choice immediately confronts us with the central conundrum: the choice should minimize the expected cost of G, but that cost depends not only on the initial choice and the costs of G1 and G2, but also on subsequent revisions of the initial choice prompted by reconsideration; i.e., it depends on the very strategy we are trying to formulate.

Our solution is approximate, and involves two stages of refinement. In the first stage, we assume that the subgoals G1 and G2 are one-step processes requiring fixed investments of effort if tried at all. (We could assume different success costs for success and failure of each goal, but the result turns out to be the same.) Under that assumption, the behaviour of an optimal strategy subsequent to the initial choice is indeed predictable: if the chosen subgoal, say G1, succeeds, then G will have been achieved at cost t_1 (on the average); if it fails, then G2 will be tried and its cost t_2 will add to that of G1, with success or failure of G at that point. Therefore the expected cost of G when starting with G1 is $E_1 = t_1 + (1-p_1)t_2$. Similarly, the expected cost of G when starting with G2 is $E_2 = t_2 + (1-p_2)t_1$. Hence, we should initially choose G1 iff $E_1 \leq E_2$, or equivalently, iff $t_1/p_1 \leq t_2/p_2$.

The ratio t/p seems like a rather natural measure of the "unworthiness" of a goal, since it is high to the extent that the goal's expected cost is high and its success probability low (cf. Simon & Kadane, 1975). Nevertheless, choices based on this ratio can be seriously in error. This is evident from the two-box monkey-and-bananas problem described above. Assuming that the probability of success with either box is 1, then a monkey choosing on the basis of t_1/p_1 and t_2/p_2 (with $p_1 = p_2 = 1$) would opt for the more remote box (which has lower expected cost), even though a small investment in the seemingly more expensive box (e.g., a trial push) might reveal it to be cheaper by far. Of course, the error stems from the assumption that pursuit of a subgoal cannot be interrupted until

it succeeds or fails.

Examples such as this one motivate the second stage of refinement in the anticipation of the strategy's behaviour and hence in the expected cost calculation and the resultant choice criterion. The (t, p) characterization of a goal is inadequate for the refinement: we need a characterization which reflects not only t and p but also the anticipated changes in t and p after investment of a certain initial effort s in the expansion of the goal. The anticipated changes are given by a probability distribution $f(t', p')$ over the possible cost and probability estimates (t', p') that may result after investment of s . The mean of the distribution is $(t-s, p)$. After careful consideration of computational implications, we have chosen to approximate any such bivariate distribution by a set of at most five piecewise linear conditional densities at fixed p' , (e.g., at $p' = .1, .3, .5, .7, .9$), mixed with at most five probability mass points. The earlier simplistic assumption that the choice between G_1 and G_2 is not reconsidered until the chosen goal succeeds or fails is now replaced by the following 'one level look-ahead' view of the selection strategy's behaviour subsequent to the initial choice. If the initial choice is G_1 , then effort s_1 is invested in its expansion, whereupon (t_1, p_1) are revised to (t_1', p_1') . Then the simplistic criterion $t_1'/p_1' \leq t_2/p_2$ is used to decide whether to continue with G_1 or switch to G_2 , and the resulting choice is not reconsidered unless and until the chosen goal fails. With this refinement, the expected cost of G when starting with G_1 is found to be

$$E_1 = s_1 + P\{U\}\{E\{t_1' | U\} + (1-E\{p_1' | U\})t_2\} + (1-P\{U\})\{t_2 + (1-p_2)E\{t_1' | \text{not } U\}\},$$

where U is the event that $t_1'/p_1' \leq t_2/p_2$; i.e., U is the region in t', p' -space which lies to the left and above a radial line from the origin with constant slope t_2/p_2 . The formula for E_2 is analogous. E_1 and E_2 are readily estimated from the approximations to the distributions $f_1(t_1', p_1')$ and $f_2(t_2', p_2')$, and so the revised choice criterion $E_1 \leq E_2$ is readily implemented.

The required generalization to n -ary OR nodes can be accomplished by computing the cost E_1 of starting with goal G_i and after effort s_i either completing G_i or switching to try the alternatives in order of non-decreasing unworthiness t_i/p_i . The formula for E_i

turns out to be exactly analogous to that for E_1 :

$$E_i = s_i + P\{I\}\{E\{t_i' | I\} + (1-E\{p_i' | I\})C_i\} + (1-P\{I\})\{C_i + q_i E\{t_i' | \text{not } I\}\},$$

where I is the event that $t_i'/p_i' \leq C_i/(1-q_i)$, C_i is the cost of trying the remaining goals in order of non-decreasing t/p , and q_i is the probability of all remaining goals failing. (e.g., if $G_i = G_2$ and the goal ordering is G_1, G_2, G_3, G_4 , then $C_2 = t_1 + (1-p_1)t_3 + (1-p_1)(1-p_3)t_4$, and $q_2 = (1-p_1)(1-p_3)(1-p_4)$.) Because of shared terms among the E_i , the decision cost increases only linearly with the number of subgoals.

AND nodes, in a planning context, are somewhat less troublesome than OR nodes. If the subgoals constitute the steps of a plan to achieve the parent goal, then it is often reasonable to expand the subgoals in a left-to-right depth-first manner (as STRIPS does -- see Fikes & Nilsson, 1971), since elaboration of early parts of a plan helps to define the initial conditions for later subgoals. However, in a genuinely hierarchical planning system (i.e., one which is capable of representing plans at various levels of abstraction or conversely, levels of detail), such a strategy would tend to overdevelop early parts of a plan, oblivious to manifest inadequacies in its later parts. One heuristic remedy, implemented by Sacerdoti (1974) in ABSTRIPS and Siklossy & Dreussi (1973) in LAWALY, is to arrange for complete elaboration of a plan at one level of abstraction before allowing its elaboration at the next level. This strategy has the converse flaw of trying to anticipate in detail the remote future consequences of present actions; the more uncertain the planning environment, the more serious this flaw is apt to be.

We have no general improvement on such heuristic methods for controlling the elaboration of strongly interdependent subgoal sequences. (Our STRIPS-world simulations rely on left-to-right expansion of AND nodes.) For conjunctions of more or less independent subgoals, however, the dual of our method for disjunctions is available.

Again, we begin by considering the case in which two goals G_1 and G_2 are uninterruptable, single-step processes. If G_1 is tried first, then the expected cost of the superordinate AND node G is $t_1 + p_1.t_2$, since G_2 will be tried only if G_1 succeeds. Similarly the expected cost of starting with G_2 is $t_2 + p_2.t_1$.

The criterion for choosing G1 is then easily seen to be

$$t1/(1-p1) \leq t2/(1-p2).$$

By an argument completely analogous to that for OR nodes, we then move to the second stage of refinement, obtaining the expected cost of beginning with G1 (and allowing for reconsideration)

$$E1 = s1 + P[V]\{E[t1'|V] + E[p1'|V]t2\} + (1 - P[V])\{t2 + p2.E[t1'|not V]\},$$

where V is the event that $t1'/(1-p1') \leq t2'/(1-p2')$.

4. Backing up

The remaining problem is the specification of backed-up characterizations at a goal as a function of the characterizations of its subgoals. Under the assumption of goal independence, the backed-up probability is

$$p = 1 - (1-p1)(1-p2)...(1-pn)$$

at an OR node and

$$p = p1.p2...pn$$

at an AND node. The backed-up expected cost is just the expected cost E1 corresponding to the chosen subgoal, as specified above. (The generalization to n subgoals can be obtained in the conjunctive case just as in the disjunctive case.) The backed-up initial step size, likewise, is that of a chosen subgoal.

However, we also need to back up $f(t',p')$ distributions and this poses serious problems. The problems come as something of a surprise, since the assumed behaviour of the goal selection strategy subsequent to the initial choice of a subgoal quite unambiguously determines the distribution of the parent goal. In particular, suppose that the subgoal G1 is initially chosen in favour of alternative G2. Consider t' , the expected remaining cost of the OR given that the initial step of G1 has been taken (at cost $s1$), and we have values for $t1$ and $p1$. If $t1'/p1' \leq t2/p2$ (event U), we continue with G1 and $t' = t1' + (1-p1')t2$. If not, we switch to G2 and $t' = t2 + (1-p2)t2'$. Similarly the probability of the OR is $p' = 1 - (1-p1')(1-p2)$. These equations determine two linear transformations of the jointly distributed random variables $(t1',p1')$, where the first transformation applies in the region U above and to the left of

the radial line with slope $t2/p2$, and the other applies below and to the right of that line. The distribution of (t',p') can be computed accordingly.

The trouble with such backed-up distributions is evident from the example shown in Fig. 1. Here the




Fig. 1. Computing the backed-up distribution for an OR on the basis of expected cost calculations that do not separate the cost components corresponding to success and failure of the chosen subgoal G1. The solid prism on the right is the transform of that on the left, and similarly for the prisms drawn in broken lines.

initial step for G1 is free ($s1=0$), the distribution $f1(t1',p1')$ is uniform on the square, and the characteristics of the alternative are $(t2,p2) = (.5,.5)$. Since both linear transformations involve a rightward shift, the resulting distribution assigns a significantly smaller probability to the event that the OR will be solved quickly than did the distribution for G1 alone (e.g., the probability of successful termination at cost $\leq .1$ drops from .05 to less than .01). Yet we know that making available an alternative to a goal cannot have a deleterious effect.

One diagnosis of the trouble is that by lumping together the costs of success and failure, we have ended up adding costs associated with failure to $t1'$ (namely the cost of proceeding with alternative G2), regardless of whether expenditure of $t1'$ led to success or failure. A possible remedy is the replacement of the distributions over (cost, probability) pairs by distributions over (cost of success, cost of failure, probability) triples, so that the two kinds of cost can be separately backed up. However, we are hesitant to go from an already complex two-dimensional parameterization to a three-dimensional one.

A less drastic, albeit somewhat ad hoc, remedy is to treat the expected costs of success and failure as identical, and thus to view the two-dimensional distribution as implicitly supplying both. We can then avoid degradation of the preferred goal G1 by alternative G2 as follows. For the case $t1'/p1' \leq t2/p2$ (where we continue with G1) we apply separate transformations

corresponding to success and failure of G_1 . We have $t' = t_1'$ if G_1 succeeds and $t' = t_1' + t_2$ if G_1 fails. We recombine the two resulting densities additively, weighting them by the success and failure probabilities p_1' and $(1-p_1')$ respectively. The probability transformation is the same as before, and so is the cost transformation in the region $t_1'/p_1' > t_2/p_2$. The effect of the transformations is illustrated in Fig. 2, for the previously considered




Fig. 2. Backed-up distribution based on splitting the cost of beginning with G_1 into a success part (see broken lines for corresponding region) and a failure part (see dotted lines for corresponding region). Only the part of the transformed distribution corresponding to continuation with G_1 after the initial step as is shown.

example. Note that the probability mass lying in the vicinity of $t_1' = 0$ has been preserved by the new transformation.

The method we suggest for n -ary OR nodes is to treat them as nested sets of $(n-1)$ binary ORs. We find the backed-up distribution for the chosen subgoal combined with the first-ranked remaining goal (with rank according to t_i/p_i), then combine the result with the distribution of the second-ranked remaining goal, and so on.

Analogous methods can be formulated for AND nodes, where the subgoals are independent. However, we have been more interested in the case where the subgoals constitute the (usually strongly interdependent) steps of a plan. Here a satisfactory way to back up characterizations seems hard to determine. A workable, but quite ad hoc method is as follows. If the planning schedule is $(G_1 \text{ THEN } G_2 \text{ THEN } \dots \text{ THEN } G_n)$, then the backed-up $f(t', p')$ is $f_1(t', p')$ translated by $(p_1(t_2 + (p_2(t_3 + \dots + p_n t_n) \dots)))$ in the p' -direction and with scale factor $p_2.p_3 \dots p_n$ applied to p' .

For both AND and OR nodes, the backed-up initial step size s is that of the chosen subgoal. Multi-step processes can be broken down hierarchically into two-step processes; a theoretical analysis indicates that the dividing point should be chosen in such a way

that the standard deviation of remaining cost is greater than 2.5 times first stage cost.

A final generalization we have studied concerns goals which may not terminate, i.e., whose failure may never become apparent. We find that such goals can be treated as if they terminated after expenditure of total effort T at the latest, where T is chosen to equal twice the expected cost when the maximal investment is T . This cut-off minimizes the effective "unworthiness" t/p .

5. Evaluation of the strategy

How useful is the overall strategy? The manipulation of approximate probability distributions may seem a little forbidding, but it should be kept in mind that in actual use these distributions can be very simple; e.g., they may be two-point or even one-point distributions. More complicated distributions are justified only in very complex planning domains.

We have several indications that the strategy is well-behaved in its decision-making. Naturally, it overcomes the sort of difficulty illustrated with the two-box monkey-and-bananas problem, as the reader can verify. (The box next to the monkey has a low initial-step cost and is characterized by a distribution with two probability mass points (t_1', p_1') and (t_1'', p_1'') , where t_1' is low, t_1'' is high, and $p_1' = p_1'' = 1$; the more remote box has cost intermediate between t_1' and t_1'' and probability 1. The proximate box will be tried first in most cases, even if its expected cost exceeds that of the more remote one.)

In a version of the multi-box monkey-and-bananas problem, recast as already explained as the problem of locating and buying a low-cost item such as a snowshovel, the strategy has actually been proved optimal (cf., Feldman and Sproull's 'wheel-less student' problem). The problem assumes that the purchaser lives at the centre of a large circular city and possesses a newspaper containing numerous ads offering snowshovels for sale. He has set aside sufficient cash, and valuing his time, wishes to obtain a shovel as quickly as possible. To locate a shovel, he must find the next shovel-ad and phone the vendor, an operation requiring unit time (cost). At any time, he may stop locating shovels and go buy the nearest one so far at cost az , where z is the distance to the vendor divided by the

radius of the city. Assuming a uniform distribution of snowshovels for sale over the city, it can be shown that the minimal expected total cost is assured by locating shovels until one is found within distance T times the radius of the city, with $T^3 = 3/a$. Precisely the same stopping criterion is obtained with the one level look-ahead strategy, with alternative subgoals "buy the nearest shovel so far" and "locate additional shovels before buying one", where the initial step size for the latter goal is $s_2 = 1$ and $f_2(t_2', 1)$ is based on the assumption of uniformly distributed snowshovels. Thus, the one level look-ahead strategy can yield optimal planning decisions.

The strategy has also been tested by hand-simulation in a STRIPS-type world, consisting of a set of 5 connected rooms, a mobile robot, 4 pushable boxes in various rooms, and a lightswitch requiring a box to reach. The robot has 8 operators, with pre- and postconditions in the quantifier-free form employed by LAWALY (Siklossy & Dreussi, 1973). Two planning problems were simulated in detail, one requiring any two boxes to be brought next to each other, and the other requiring the lightswitch (not located in the same room as the robot) to be turned on. For computing cost estimates, unit cost was assigned to performing a match of a literal against the world model and to attaching a node in the planning tree, and a cost "a" to executing an operator. All goals generated were assumed to be achievable (probability 1), allowing the required distributions over cost and probability to be collapsed into simple discrete distributions over cost alone. For example, the distribution pertaining to the goal of being in a specified room is approximated by a two-point probability mass function, obtained by assuming that for an initial room R' , target room R and any room R connected to R' , there is a 50% chance that R is "one door closer" to R' than R (in terms of the minimum number of doors that have to be passed through). For simplicity, goal conjunctions were scheduled left to right.

The simulations, which produce about 100 goal nodes each, nicely illustrate first, how the control strategy abandons unpromising plans before fully elaborating them (unlike failure-driven backtrack search) and second, how a goal of higher expected cost may be selected over a goal of lower expected cost because it offers some chance of a very cheap solution.

In the future, it would be useful to implement the strategy and test it on more complex examples, particularly examples with non-unit goal probabilities. Also, the problems encountered in backing up distributions call for further research, perhaps in the direction of separate characterizations for success and failure. Another difficult issue calling for further research is that of goal interdependence; two types of interdependence are of particular interest, namely the dependence of later subgoals in a plan on earlier subgoals, and the possible dependence of many subgoals in many competing plans on a common information-gathering action, such as visual search or a feasibility test (e.g., checking a box for pushability).

References.

- Feldman, J. A., and Sproull, R. F. (1974). Decision theory and artificial intelligence II: The hungry monkey. Tech. Rep. No. 2, Computer Science Dept., Univ. of Rochester, Rochester, NY. Also Cognitive Science 1, 158-192 (1977).
- Fikes, R. E., and Nilsson, N. J. (1971). STRIPS: A new approach to the application of theorem proving to problem solving. AI 2, 189-208.
- Johnson, D. T. (1977). Decision theory and automatic planning. Tech. Rep. TR77-9, Dept. of Computing Science, Univ. of Alberta, Edmonton, Alberta.
- Nilsson, N. J. (1980). Principles of Artificial Intelligence. Tioga Publ. Co., Palo Alto, CA.
- Sacerdoti, E. D. (1974). Planning in a hierarchy of abstraction spaces. Artificial Intelligence 5, 115-135.
- Siklossy, L., and Dreussi, J. (1973). An efficient robot planner which generates its own procedures. Proc. IJCAI 3, 423-430.
- Simon, H. A., and Kadane, J. B. (1975). Optimal problem-solving search: all-or-none solutions. AI 6, 235-247.
- Slagle, J. R., and Bursky, P. (1968). Multipurpose theorem-proving heuristic program. JACM 15, 85-99.