

Ordering Relations in Human and Machine Planning¹

Lee Spector, Mary Jo Rattermann, and Kristen Prentice

School of Communications and Cognitive Science
Hampshire College, Amherst, MA 01002
{lspector, mrattermann, kprentice}@hamp.hampshire.edu

Abstract

Analytical results from AI planning research provide the motivation for this experimental study of ordering relationships in human planning. We examine timings of humans performing specific tasks from the AI planning literature and present evidence that normal human planners, like “state of the art” AI planning systems, use partial-order plan representations. We also describe ongoing experiments that are designed to shed light on the plan representations used by children and by adults with planning deficits due to brain damage. Several points of interest for collaboration between AI scientists and neuropsychologists are noted, as are impacts that we feel this research may have on future work in AI planning.

Introduction

Recent analytical studies have mapped a rich territory of relations among AI planning algorithms, AI planning domains, and computational complexity classes (Chapman 1987, Bylander 1991, Chenoweth 1991, Gupta & Nau 1991, Erol et al. 1991, Minton et al. 1992, Barrett & Weld 1993). The work of Barrett and Weld, in particular, shows that planning algorithms based on partially-ordered plan representations have advantages over total-order planners when applied to problems of certain sorts. One of their results is that problems in the artificial domain called D¹S¹ are easy for partial-order planners but hard for total-order planners. Partial-order planners exhibit linear growth rates for planning time as problem size increases in this domain, while total-order planners exhibit exponential growth rates.²

This result motivates the present study of human planning. We know that partial-order planners have specific advantages over total-order planners in AI systems; by looking for signs of such advantages in human systems, we can draw tentative inferences about the representations and algorithms used in human planning. The lessons learned about human

planning can then be re-applied to AI systems, which still lag behind human competence in several areas.

We presented human subjects with tasks from the D¹S¹ domain and noted the time spent planning, as well as the overall time needed to complete the task. As in AI planning systems, linear growth rates in human performance are suggestive of partial-order planning, while exponential growth rates are suggestive of total-order planning. An additional goal of this work is to shed light on the role of the frontal lobe of the brain in the ability of humans to perform complex planning tasks. To this end, we plan to compare the performance of normal adults, adults with frontal lobe lesions, and children at various stages of frontal lobe development.

The remainder of the paper is organized as follows: after a brief refresher on partial-order planning we summarize the relevant complexity results of Barrett and Weld and discuss the neuropsychological background of our experiment. We then describe the Chores software with which the experiment was conducted, along with the details of the D¹S¹ task. This is followed by a summary of our results to date and a discussion of implications and directions for future work.

Partial-order Planning

Most AI planning systems since STRIPS (Fikes & Nilsson 1971) have represented actions as *operator schemata*. Operator schemata describe the conditions under which each action can be performed and the effects that the performance of each action will have on the world. STRIPS-style operator schemata are composed of *precondition* lists that represent the conditions under which the actions can be performed, and pairs of lists called *add* and *delete* lists that represent effects. A *plan* to achieve some given goal from some given initial state is a list of operator schemata, along with assignments for all variables in the schemata, such that the corresponding sequence of actions will appropriately transform the world.

Early planners such as STRIPS worked by constructing sequential (linear, total-ordered) *partial plans* which achieved subsets of the goals, and which could be augmented, manipulated and combined to produce correct final plans. Given a list of goals to achieve, and a system that supports linear partial plans, the natural strategy is to produce plans to achieve each of the goals and to concatenate the resulting partial plans. This strategy can be applied re-

¹The first author acknowledges the support of the Dorothy and Jerome Lemelson National Program in Invention, Innovation, and Creativity.

²These results have only been proven for the particular planning algorithms in the Barrett and Weld study, but we believe them to be indicative of the advantages of partial-order planning more generally. (Minton et al. 1992) provides some support for this belief.

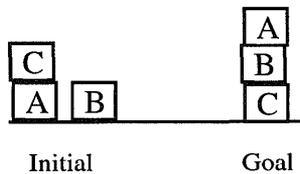


Figure 1. The Sussman Anomaly

cursively to any subgoals introduced in trying to achieve any of the goals. But using linear partial plans in this way leads to a difficulty that can be illustrated with a blocks world problem known as the Sussman Anomaly: Given three blocks labeled A, B, and C, and an initial state in which C is on A, and in which A and B are both on the table, construct a plan for building a tower with A on B and with B on C (Figure 1). The difficulty is that no linear combination of the obvious partial plans for achieving $\text{On}(A, B)$ and $\text{On}(B, C)$ will yield a correct plan for achieving their conjunction. If $\text{On}(A, B)$ is achieved first (by putting C on the table and then putting A on B) then the new {A, B} stack would have to be immediately dismantled in order to achieve $\text{On}(B, C)$. (Only one block can be lifted at a time.) On the other hand, if $\text{On}(B, C)$ is achieved first (by putting B on C) then once again the new stack would have to be dismantled in order to achieve the other conjunct.

The literature contains several strategies for circumventing this problem (e.g., (Sussman 1990, Sacerdoti 1975, Waldinger 1977)). One of the most popular strategies involves the representation of partial plans as *partial-orders*, rather than as *total-orders*. This allows for the representation of partial plans in which some of the ordering decisions have been left temporarily unspecified. A partially ordered partial plan for the Sussman Anomaly might state that $\text{On}(A, B)$ and $\text{On}(B, C)$ are both to be achieved, but that the order in which their plans are to be executed has yet to be determined. Further computation may indicate that additional ordering constraints must be imposed; in the case of the Sussman Anomaly the actions into which the goals are decomposed will have to be *interleaved*. Figure 2 illustrates a solution to the Sussman Anomaly using partially ordered action sequences.³

Most “state of the art” planning systems now use partially-ordered partial plans (e.g., (Chapman 1987, Currie & Tate 1991, McAllester & Rosenblitt 1991)). Several recent studies have examined the trade-offs in total-order vs. partial-order planning; some of these studies quantify the differences in performance profiles between total-order and

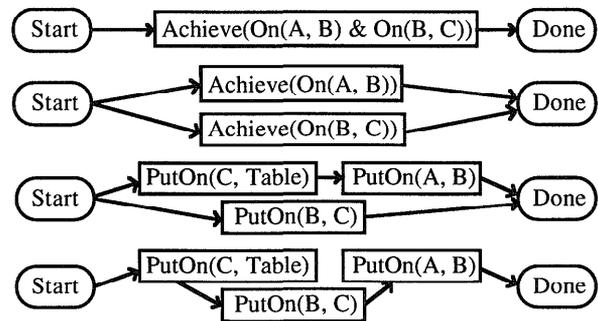


Figure 2. Solving the Sussman Anomaly with partially ordered partial plans.

partial-order planners (Minton et al. 1992, Barrett & Weld 1993). With regard to human planning, we know of no prior literature on the ordering relations in partial plans. The results from AI planning systems can, however, be used to structure experiments that provide data on the ordering relations used by humans in constructing action plans.

Barrett & Weld’s Results and Domain D’S¹

Barrett and Weld examined the performance of three planning algorithms on problems of several classes. Two of their planners are of interest here: *POCL* (for “Partial-Order, Causal-Links”) and *TOCL* (for “Total-Order, Causal-Links”). *POCL* and *TOCL* differ, to the extent that this is possible, only in the representation of ordering relations within partial plans; *POCL* uses partial-order representations while *TOCL* uses total-order representations. Both *POCL* and *TOCL* use standard STRIPS-style operator schemata. *POCL* is a variant of the *systematic* nonlinear planning algorithm used in the SNLP planning system (McAllester & Rosenblitt 1991), and *TOCL* is a modification of *POCL* in which an added “linearization” step forces all plans to be totally ordered.⁴

POCL and *TOCL* were run on large sets of problems from both “real” and artificial domains. The real domains included blocks world, transportation planning, and others, while the artificial domains were constructed to highlight the effects of specific patterns of goal interactions. The runtimes for each planner on problems in each domain were analyzed, and several interesting generalizations were noted.

In accounting for the observed data, Barrett and Weld extended Korf’s taxonomy of subgoal collections (Korf 1987) to include two new classes: *trivial serializability* and *laborious serializability*. A set of subgoals is said to be trivially serializable “if each subgoal can be solved in any order without ever violating past progress,” while a set of subgoals is said to be laboriously serializable “if there exist an inadequate percentage of orders in which the subgoals may be solved without ever violating past progress”⁵ (pp. 3–4). Barrett and Weld also produced the following result:

Proposition II Assuming that a problem’s subgoals can be achieved in constant time, the expected time to solve a problem rises linearly with the number of sub-

³See (Sacerdoti 1975, Tate 1977) for further detail.

⁴The third planner in the Barrett and Weld study (*TOPI*, for “Total-Order, Prior-Insertion”) differs from *TOCL* along a different dimension and hence is of no interest to the present study. Barrett and Weld mention that source code for all three planners is available—send mail to bug-snlp@cs.washington.edu.

⁵The full definition: “A set of n subgoals is laboriously serializable if there exists at least one serializable ordering yet at least $1/n$ of the subgoal orders can not be solved sequentially without possibly violating a previously solved subgoal.” (p. 19)

goals if the problem is trivially serializable, but rises exponentially if the problem is laboriously serializable or nonserializable. (p. 36)

It is important to note that the classification of a problem as trivially or laboriously serializable depends on the planning algorithm that is being used; different planning algorithms induce different subgoal structures, and hence it is possible for a single problem to be trivially serializable in the search space of one planner and laboriously serializable in the search space of another. The artificial domain called D'S¹ is the simplest of Barrett and Weld's domains that is trivially serializable for *POCL* but laboriously serializable for *TOCL*. In D'S¹ each operator has exactly one precondition, one item on its add list, and one item on its delete list. The precondition lists and add lists are all mutually disjoint, and each operator deletes the precondition of one of the other operators. The pattern for the construction of an n -operator D'S¹ domain is:

Operator:

Action: A_i , Preconditions: $\{I_i\}$, Add: $\{G_i\}$, Delete: $\{I_{i-1}\}$

where i ranges from 1 to n and where the delete list for A_1 is empty. The maximal solvable problems in this domain have initial conditions consisting of all of the I_i and goal conditions consisting of all the G_i . Note that the pattern induces a linear order on the solution to each maximal problem—each operator deletes the precondition to the operator that must immediately precede it in the solution plan. The results of Barrett and Weld tell us that partial-order planners are capable of finding the correct order in linear time, while total-order planners will generally require exponential time.

Neuropsychology of Planning

Neuropsychologists have long studied the role of the frontal cortex in human behavior, and in particular, the effects of damage to this area of the brain. It has long been known that patients with frontal lobe lesions will often behave inappropriately in social situations, experience radical mood swings, and display deficits in processing temporal relationships and order. These deficits appear despite seemingly normal abilities in language, perception, verbal expression, memory and attention (Grafman 1989, 1994; Robertson et al. 1991). In addition to the previously described deficits, it has also been proposed that these patients experience deficits in planning and problem solving (Grafman 1989, Shallice 1988). Unfortunately, the majority of neuropsychological models fail to explicate the precise role of the frontal cortex in cognitive processes (Grafman 1989), and also rely on evidence which is obtained from tasks not directly related to planning and problem solving. One of the goals of our experiment is to provide direct evidence for the role of the frontal cortex in plan-

ning behavior, and specifically, to show that damage to the frontal cortex will lead to deficits in the patient's ability to form and carry out complex plans.

D'S¹ in the Chores Experiment

We are performing experiments built using the NINDS/NIH Chores experiment software (Spector & Grafman 1994) to explore the ordering representations used in human problem solving. The software was designed for human planning experiments with normal subjects, with frontal lobe lesion patients, and with children at various stages of frontal lobe development.⁶ The subject interacts with a Macintosh computer that displays a map of a hypothetical city, with icons representing items that the subject is to acquire (the goal list), and icons representing items that the subject already possesses (the inventory). The subject uses the mouse and/or arrow keys to navigate around the map and to perform chores. The subject can backtrack by clicking an "Undo Previous" button—this reverses the effects of the most recently performed chore and moves the subject to the location adjacent to that chore. A time-stamped protocol that lists all of the subject's actions is produced; it can be analyzed to produce several data sets.

The Chores software can be configured to require that constraints be obeyed on the ordering of the chore completions. In our experiments we are using "resource" constraints that are specified in a manner almost identical to STRIPS operator specifications. On a separate screen called "Item Info" each chore is listed along with its relations to resources. The subject may switch between the map and Item Info screens at any time; the times of all such actions are recorded in the protocol. The Item Info screen may indicate that a particular chore "requires" one or more resources—this means that the chore cannot be completed if the resources are not currently in the inventory. The system produces a beep-sound and a time-stamped annotation in the protocol whenever a subject attempts to complete a chore without the necessary requirements. A chore may also "add" one or more resources—this means that successful completion of the chore will add the specified resources to the inventory. If the added item was a goal then it is also removed from the goal list. A chore may also "delete" one or more resources—this means that successful completion of the chore will delete the specified resources from the subject's inventory. Note that the require, add, and delete specifications for chores are strictly analogous to the precondition, add, and delete lists for STRIPS operators. This means that the human subjects will be performing tasks quite similar to those performed by AI planners that manipulate STRIPS operators.

In the present experiment we are interested in the relation between planing time and task size (number of goals). We are using the time spent viewing the Item Info screen as our principal indicator of planning time. The total time to trial completion is also of interest, but this will include time spent manipulating the user interface that may be unrelated to planning time.

⁶The general idea for the Chores software was derived from the chore planning experiments described in (Hayes-Roth & Hayes-Roth 1979).

The D'S' Series

The D'S' Series is a set of trials built using the Chores experiment software. The entire task consists of two training trials, four test trials, and four foil trials, which are randomly presented to the subject. The first of the two training trials introduces the subject to the basic structure and rules of the task, including the “requires” and “adds” functions, while the second training trial introduces the “deletes” and “undo previous” functions. The four D'S' trials are as follows: a 2-goal problem in a 2-operator D'S' domain, a 3-goal problem in a 3-operator D'S' domain, a 4-goal problem in a 4-operator D'S' domain, and a 5-goal problem in a 5-operator D'S' domain. The foils are similar to the D'S' trials, but they do not actually conform to the D'S' pattern. To minimize the effects of varying spatial reasoning abilities, all of the maps have the same simple linear layout.

Subjects and Methods

Seventeen normal adults and two frontal subjects were tested in this experiment. Normal subjects were screened for histories of motor and cognitive impairment, and any necessary corrective lenses were worn during testing. Frontal subjects were diagnosed to have cortical damage located in the frontal lobe.

Testing took place in a testing room with only the subject and experimenter present. The subject was seated in front of the computer monitor and was introduced to the experiment with the training trials. The following instructions are a summary of those given to the subjects:

- The goal of this experiment is to obtain a set of items by going to the places on this map. Each time an item is obtained it is removed from the list of objects to get and added to the inventory.
- The places and items on the list are not thematically linked. Thus, going to the “library” will not result in a book being added to the inventory, rather, a triangle will be added.
- The small black square which designates position can move forward or backward one square at a time, and a particular place on the map is entered by moving the black square on top of it. [Subjects are shown how to manipulate the square with the mouse and the arrow keys.]
- The “Item Info” screen displays the relationships between the places on the map and the items. If a place “requires” an item, entry is prohibited unless that item is in the inventory. If a place “adds” an item, then that item will be added to the inventory when that place is entered. If a place “deletes” an item, then if that item is present in the inventory it will be deleted. If the item is not in the inventory then the deletion has no effect.
- The “undo previous” button can be used to undo a previous move. This button may be pressed as many times as necessary to undo to the desired move.

Subjects were encouraged to practice with the training trials and to ask questions until they were comfortable with the rules of the task. The subjects were informed that the

trials would be timed and that they should move through them as quickly as possible, but that randomly guessing the order in which to complete the chores would not be beneficial.

The four D'S' trials and four foils were presented in a different random order to each subject. Upon the completion of each trial the subject protocol was saved for later examination.

Results

We analyzed two aspects of the subjects' performance in the Chores task: Total time to completion and time spent viewing the Item Info screen. The former is the amount of time spent both planning and executing the plan; the latter is a better measure of pure planning time, but it does not include time spent planning from memorized item information. We had predicted that both measures would reveal the hypothesized linear trend for the normal subjects, and as can be seen in Table 1 and in Figures 4 and 5, this prediction was confirmed by the normals' performance on both measures. A linear trend analysis performed on the data from the normal subjects reveals a significant linear trend for total time to completion ($F(1,48) = 31.14, p < .0001$) and a significant linear trend for time spent viewing the Item Information screen ($F(1,48) = 48.88, p < .0001$). Further support for our hypothesis is provided by analyses of variance performed on both sets of data which revealed a main effect of the number of Chores ($F(3, 48) = 10.48, p < .0001$ for time to completion, and $F(3, 48) = 16.74, p < .0001$, for time spent viewing the Item Information screen).

Due to the small number of frontal subjects, statistical analyses were not possible; however, as can be seen in Figures 4 and 5, their preliminary data is consistent with the predicted exponential function.

Discussion

Based on suggestions that planning algorithms using partially-ordered plan representations have clear advantages over total-order planners in particular AI planning domains, we set out to investigate the performance of human planning systems in the same domains. Our results suggest that our human subjects enjoy the same advantages. Using the Chores software to test subjects in D'S' domains, we found that the subjects' total time to completion and time spent viewing the Item Information screen exhibited linear growth rates as the problem size increased.

The data suggests that human planners represent the or-

Number of Chores	Time to Completion	Viewing Item Information
Two Chores	54.91	12.35
Three Chores	131.76	48.71
Four Chores	225.12	93.12
Five Chores	343.53	159.76

Table 1. Mean total time and mean time spent viewing Item Info.

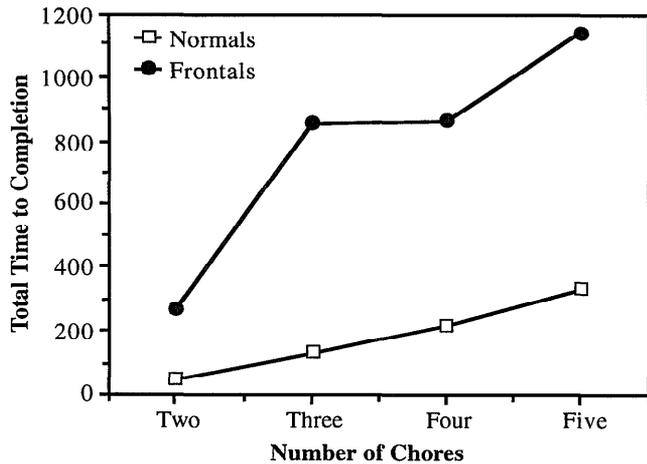


Figure 4. Mean Total Time to Completion, in seconds

dering relations in partial plans using partial-orders. The data only *shows* that human planning time seems to increase linearly with increase in task size in the D'S1 domain. The inference that this implies the use of partial-order representations relies on the assumption that the results of Barrett and Weld are generally indicative of the advantages of partial-order planning as opposed to total-order planning. Although we believe this to be the case, we note that the experiments provide data even if the results of Barrett and Weld fail to generalize. The human planners are managing to achieve linear performance *somehow*. They are either using partial-order representations or other algorithmic methods that achieve the same effect with respect to efficiency in this domain. One alternative is that the linear human performance owes to the brain's use of parallel algorithms.⁷ While such alternatives deserve further study, we currently believe that partial-order representations provide the best explanation for the data.

An additional goal of this experiment was to examine the role of frontal cortex in planning by testing patients with frontal cortex lesions. Although preliminary, our results suggest that the performance of frontal patients in this domain is more similar to that of total-order planners; specifically, their performance seemed to exhibit an exponential growth rate as problem size increased. More data from frontal patients is necessary to strengthen this claim.

Our results have implications for both the psychology and the neuropsychology of planning, as well as for the AI community. First, we have provided evidence that in a domain in which partial-order planning is advantageous, normal adults perform in a way that suggests the representation of partial-plan ordering relations as partial-orders. Second, we have preliminary evidence that suggests that the frontal cortex of the brain may play a role in this type of planning. While frontal lesion patients' difficulty with social situations and with temporal ordering is well docu-

⁷We thank an anonymous reviewer for drawing our attention to this possibility.

mented, the effect of this damage on planning is not as well studied. In our work, we have directly tested the ability of these individuals to form and carry out plans, and we believe that their deficits lead to a specific pattern of performance with analogs in the computational literature.

The utility of partial-order planning, discovered by AI planning researchers, receives further support from the evidence that normal humans appear to use partial-order plan representations. Further study may uncover additional features of the algorithmic basis of human planning, both in normal subjects and in subjects with planning deficits. These studies may provide further support for current AI techniques, but they may also provide alternative models. Given that human planners are proficient in ways that current AI models are not, the interplay between the human and machine studies should be of interest to those extending the state of the art in AI planning. The fruitfulness of this interplay relies, however, on the coherence of the cross-disciplinary dialogue. We suggest that one way to maximize the impact of psychological experimentation on AI practice is to construct the experiments within the conceptual frameworks developed in AI research.

Future Work

In our ongoing studies we are continuing our examination of the role of frontal cortex in human planning. We are currently testing more frontally lesioned patients, and we have adapted the procedure to make it appropriate for testing 7–11-year-old children. Based on research suggesting that the frontal cortex does not fully develop until late in childhood, we expect to find that the younger children display the same exponential performance that our frontal subjects exhibit. This work will provide further evidence for the role of the frontal cortex in human planning behavior.

Although the present study focuses on the representation of ordering relations in partial plans, the same framework can be applied to study other aspects of planning, and other aspects of cognition more generally, across the human/machine frontier. AI research often yields precise, quantitative

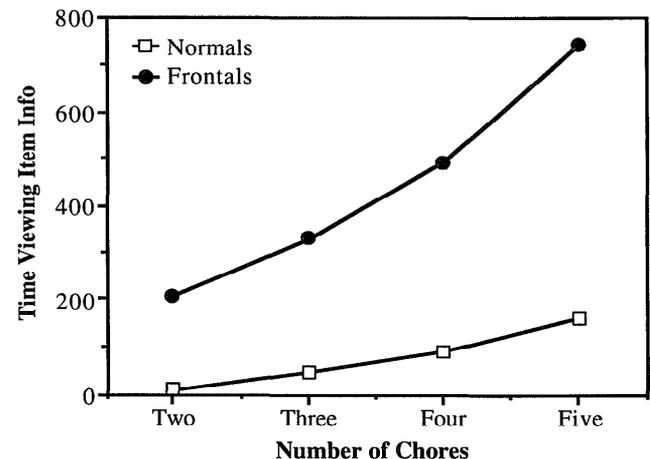


Figure 5. Mean Time Viewing Item Information, in seconds.

results that relate computational structures to aspects of performance. By using these results as guides, we can gather data on human cognition that will allow us to make inferences about human computational structures. And of course there is always the possibility that we will be surprised; that we will find that the human computational structures were other than we had expected. Since the concepts that underlie these investigations will be taken from the computational literature, there should be clear paths by which one could apply the lessons of the surprises in improving the state-of-the-art in AI.

Conclusions

In conclusion, we have found that the results of analytical work in AI planning can be used to investigate human planning. Specifically, we have evidence that suggests that normal human planners use partial-order representations for partial plans, as do most modern AI planning systems. We have further preliminary evidence that suggests that damaged human planning systems use methods akin to those used in less efficient AI systems; specifically, we believe that adults with frontal lobe lesions will be shown to be using planning methods similar to those employed in total-order planners. We believe that the parallels between the human and machine cases are instructive, and that they may lead to further developments in both human and machine studies.

Acknowledgments

Jordan Grafman funded and guided the development of the Chores software, and, along with Jim Hendler, was instrumental in launching the larger interdisciplinary effort of which this work is a part. Paula Koseff ran several trials of the experiment on subjects at NIH under Grafman's supervision. Chris Chase and Neil Stillings helped in analyzing the results. Jason Juneau worked on the software and is part of the ongoing experimental team.

References

- Barrett, A.; and Weld, D.S. 1993. Partial-Order Planning: Evaluating Possible Efficiency Gains. Technical Report 92-05-01 Expanded Version, Dept. of Computer Science and Engineering, U. of Washington, Seattle.
- Bylander, T. 1991. Complexity Results for Planning. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence, IJCAI-91*, 274–279.
- Chapman, D. 1987. Planning for Conjunctive Goals. *Artificial Intelligence* 32: 333–377.
- Chenoweth, S.V. 1991. On the NP-Hardness of Blocks World. In *Proceedings of the Ninth National Conference on Artificial Intelligence, AAAI-91*, 623–628.
- Currie, K.; and Tate, A. 1991. O-plan: the open planning architecture. *Artificial Intelligence* 52: 49–86.
- Erol, K.; Nau, D.S.; and Subrahmanian, V.S. 1991. Complexity, Decidability and Undecidability Results for Domain-Independent Planning. CS-TR-2797, Dept. of Computer Science, U. of Maryland.
- Fikes, R.E.; and Nilsson, N. 1971. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. In *Artificial Intelligence* 2, 189–208.
- Grafman, J. 1989. Plans, Actions, and Mental Sets: Managerial Knowledge Units in the Frontal Lobes. In *Integrating Theory and Practice in Clinical Neuropsychology*, 93–138. New Jersey: Lawrence Erlbaum Publishers.
- Grafman, J. 1994. Alternative Frameworks for the Conceptualization of Prefrontal Lobe Functions. In *Handbook of Neuropsychology*, Vol. 9. Boller, F.; and Grafman, J., eds. Elsevier Science Publishers.
- Gupta, N.; and Nau, D.S. 1991. Complexity Results for Blocks-World Planning. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence, IJCAI-91*, 629–633.
- Hayes-Roth, B.; Hayes-Roth, F. 1979. A Cognitive Model Planning. *Cognitive Science* 3: 275–310.
- Korf, R.E. 1987. Planning as Search: A Quantitative Approach. *Artificial Intelligence* 33: 65–88.
- McAllester, D.; and Rosenblitt, D. 1991. Systematic Non-linear Planning. In *Proceedings of the Ninth National Conference on Artificial Intelligence, AAAI-91*, 634–639.
- Minton, S.; Drummond, M.; Bresina, J.L.; and Philips, A.J. 1992. Total Order vs. Partial Order Planning: Factors Influencing Performance. In *KR '92. Principles of Knowledge Representation and Reasoning — Proceedings of the Third International Conference*, 83–92. San Mateo, CA: Morgan Kaufmann Publishers.
- Robertson, L.C.; Lamb, M.; and Knight, T. 1991. Normal Global-Local Analysis in Patients with Dorsolateral Frontal Lobe Lesions. *Neuropsychologia* 29, 959–967.
- Sacerdoti, E.D. 1975. The Nonlinear Nature of Plans. In *Advance Papers of the Fourth International Joint Conference on Artificial Intelligence, IJCAI-75*, 206–214.
- Shallice, T. 1988. *From Neuropsychology to Mental Structure*, New York: Cambridge University Press.
- Spector, L.; and Grafman, J. 1994. Planning, Neuropsychology, and Artificial Intelligence: Cross-Fertilization. In *Handbook of Neuropsychology*, Vol. 9, Boller, F.; and Grafman, J., eds. 377–392. Amsterdam: Elsevier Science Publishers.
- Sussman, G.J. 1990. The Virtuous Nature of Bugs. In *Readings in Planning*. Allen, J.; Hendler, J.; and Tate, A., eds. 111–117. San Mateo, CA: Morgan Kaufmann Publishers, Inc.
- Tate, A. 1977. Generating Project Networks. In *Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI-77*, 888–893.
- Waldinger, R. 1977. Achieving Several Goals Simultaneously. *Machine Intelligence* 8 (1977).