


RL-TR-91-431
Final Technical Report
December 1991


AD-A253 102



STRATEGIES AND COST IN PLAN EXECUTION

CALSPAN-UB Research Center

James F. Allen, Steve Hanks, Josh Tenenberg,
Piero B. Bonissone, Peter C. Halverson, Marc Vilian



APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

92-20240



Rome Laboratory
Air Force Systems Command
Griffiss Air Force Base, NY 13441-5700

92 7 27 156

This report has been reviewed by the Rome Laboratory Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RL-TR-91-431 has been reviewed and is approved for publication.

APPROVED:



LOUIS J. HOEBEL
Project Engineer

FOR THE COMMANDER:



JOHN A. GRANIERO
Chief Scientist for C3

If your address has changed or if you wish to be removed from the Rome Laboratory mailing list, or if the addressee is no longer employed by your organization, please notify RL (C3CA), Griffiss AFB NY 13441-5700. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE December 1991		3. REPORT TYPE AND DATES COVERED Final Jan 91 - Mar 91	
4. TITLE AND SUBTITLE STRATEGIES AND COST IN PLAN EXECUTION				5. FUNDING NUMBERS C - F30602-88-D-0026 Task 45 PE - 62702F PR - 5581 TA - 27 WU - PH	
6. AUTHOR(S) James F. Allen, Steve Hanks, Josh Tenenberg, Piero B. Bonissone, Peter C. Halverson, Marc Vilian					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Calspan-UB Research Center 4455 Genessee Street Buffalo NY 14225				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Rome Laboratory (C3CA) Griffiss AFB NY 13441-5700				10. SPONSORING/MONITORING AGENCY REPORT NUMBER RL-TR-91-431	
11. SUPPLEMENTARY NOTES Rome Laboratory Project Engineer: Louis J. Hoebe1/C3CA(315) 330-3655					
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This report documents the seminar and workshop on Resource Bounded Planning. The seminar consisted of a 10 week discussion group focusing on the foundations and historical development of the planning field and the state of the art as it applies to applications of interest to the Air Force. The workshop provided an open forum for invited researchers to present their "work in progress" to a research and application oriented audience. Position statements by workshop presenters and a post workshop paper are included.					
14. SUBJECT TERMS Planning, Plan Execution, Resource Constraints, Execution Costs				15. NUMBER OF PAGES 32	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT SAR		

Costs and Strategies in Plan Execution

Overview of Report

This report documents the seminar and workshop on computational (cost) issues in planning and plan execution held at Rome Laboratory in the Spring of 1991. The motivation for such was in large part due to Rome Laboratory's leading role in the DARPA/RL Crisis Planning Initiative. The seminar discussion series helped provide background necessary for continued leadership in the management of the initiative and also for active participation in the workshop. The workshop brought together researchers with various views on the notion of resource bounded planning. The discussions brought out that although there is wide agreement on the importance and nature of the problem there is no consensus on an approach to the solution itself. It was offered and agreed that some cohesive effort is needed towards formulating an approach to a solution. A first step could a wide spectrum research effort and perhaps more workshops but with a wider participation.

This report then contains a capsule summary of the entire effort and is followed by position statements offered by the presenters at the workshop. This collection includes a post workshop offering by Steve Hanks detailing some of the points of contention and agreement.

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

Strategies and Cost in Plan Execution

TASK No. B-1-3366

Final Report

Seminar on Planning

The 12-week seminar provided a highly interactive setting in which participants gained access to the planning research literature. The seminar began with the foundations of the field and its historical development and lead to an examination of the state of the art in resource bounded planning. In order to foster interaction, participants were encouraged to comment and ask questions at any time during meetings. Time was allocated for discussing topics of direct interest to participants, particularly their research projects and projects they manage. The seminar provided an informal forum in which to present this material and obtain feedback.

The seminar examined the state of the art in planning under resource limitation and uncertainty, beginning with an examination of foundational concepts and the paradigmatic problems that have historically driven developments in the field. Foundational concepts include those involved in knowledge representation, computational complexity, probability and decision theory. Historical problems include the frame problem, the qualification problem and the problems of representing time and actions. Interest now in practical applications such as human aids for large-scale planning problems has focused research on resource limitations and uncertainty. The seminar surveyed these advanced topics covering in particular reactive planning and formalisms that directly model uncertainty. The references that were covered are listed below.

Workshop on Resource-Bounded Planning

We held a small informal 2-day workshop to discuss issues in Resource-Bounded Planning. The workshop brought together a small number of researchers that are addressing different aspects of the problem and provided ample time for an in-depth presentation of each speaker's recent work and follow-up discussion.

The topic of resource boundedness is central to much current planning research. Recognizing resource bounds makes explicit the limitations faced by real agents. A real agent interacting with its environment has to deal with both complexity and uncertainty. The complexity of interpreting sensors and generating and executing plans must not exceed the computational resources available to the agent. The agent must also adopt problem solving strategies that accommodate the uncertainty in its sensory measurements and the uncertainty associated with the effects of physical actions. An agent is unlikely to have the leisure to adopt such cautious or redundant strategies as to eliminate this uncertainty. By focusing directly on an agent's resource limits, it may be possible to bridge the gap between formal models and practical applications of automatic planning.

The workshop participants included Center Program Managers at the Rome Laboratory. The workshop was part of a seminar entitled "Strategies and Costs of Plan Execution" attended by managers. The goal of the seminar was to provide an overview of the state of the art in planning research. As an introduction, the seminar started with a review of traditional problems, such as the frame problem, that have driven planning research.

The 2-day workshop consisted of hour-long presentations by invited speakers followed by a half-hour discussion period. In addition, a panel discussion was scheduled to obtain a synthesis of the speakers' different approaches. The invited speakers were:

James Allen, University of Rochester
Josh Tenenbergs, University of Rochester
Piero Bonissone, General Electric
Steve Hanks, University of Washington
Marc Vilain, MITRE Corporation

Below are summaries of their presentations.

Resource-Bounded Planning

James Allen
Dept. of Computer Science
University of Rochester
Rochester, NY 14627-0226

In a very real sense, planning is always resource-bounded, in that we are always hampered by lack of knowledge about the world, and cannot create plans that can cover every contingency. So the abstract notion of proving a plan absolutely correct is not realizable in practice. Of the techniques that are found in current planning research, three features seem to have the most relevance to this problem, namely abstraction, time, and probabilistic models. These have important implications as to the knowledge representation needed to represent the world, and plans in the world.

Steve Hanks
Department of Computer Science & Engineering, FR-35
University of Washington
Seattle, WA 98195

I see two sets of issues, one involving representation and one control. The representation problem arises in trying to express the tradeoff in the first place. One must have some notion of the consequences of making these plan/act decisions: what kinds of gains do we expect from thinking, what kinds of losses do we expect if we fail to act. Probability and utility models seem right for this task. Previous and current work of mine involves trying to build these models: how to maintain a probabilistic model of the world, and how to build a utility model appropriate for a particular set of goals. Decision theory provides a formally satisfying model for making planning decisions, though decision theory's exact role in the planning process, and more particularly how to apply the theory in a computationally feasible manner, remain as open issues.

Searching for Relevance

Josh Tenenber
Dept. of Computer Science
University of Rochester
Rochester, NY 14627-0226

It is well known that expressive planning representations, such as those that are logic based, are expensive to compute with. Typically, only weak search methods can be brought to bear to find solutions to problems. Even with less expressive representations, such as table-driven, reactive systems, there can be considerable expense associated with building tables that are both robust, and embody rational decision strategies. In both of these cases, increased knowledge about the problem domain, rather than improving problem solving efficiency, usually results in performance degradation. The problem is that these systems are unable to distinguish between relevant and irrelevant knowledge for a given task. An important theme that emerges, then, is how we might automate the agent's search for relevance in order to improve its performance.

Resource-bounded Planning


Marc Vilain
MITRE Corporation
Mail Stop A114
Burlington Rd.
Bedford, MA 01730

Among various approaches to resource-bounded planning, some of the most promising -- and most provocative -- are based on notions of machine learning. These approaches strive to improve a planner's performance over time by deriving control knowledge from past attempts at solving problems. By analyzing what worked and didn't work in solving a particular problem, these techniques can effectively learn shortcuts towards solving future related problems.

Piero P. Bonissone and Peter C. Halverson

Artificial Intelligence Laboratory
GE Corporate Research and Development Center
Schenectady, NY

Our work on resource-bounded reasoning processes has concentrated on deductive reasoning systems, identifying issues associated with making inferences in the presence of uncertain and incomplete input data, unpredictable event sequences, and critical time and resource constraints. The reasoning systems used in our architecture have been specifically designed to deal with time-constrained situations described by imprecise, uncertain information.



Workshop on Resource-bounded Planning

Rome Laboratories
Griffiss Air Force Base

March 11-12, 1991

Thoughts on Resource-Bounded Planning

James Allen

March 6, 1991

In a very real sense, planning is always resource-bounded, in that we are always hampered by lack of knowledge about the world, and cannot create plans that can cover every contingency. So the abstract notion of proving a plan absolutely correct is not realizable in practice. Of the techniques that are found in current planning research, three features seem to have the most relevance to this problem, namely abstraction, time, and probabilistic models. These have important implications as to the knowledge representation needed to represent the world, and plans in the world.

Abstraction is essential, as the actual details for any actual plan execution are impossible to work out in advance in all but the simplest situations. But at a suitably abstract level, one might formulate a plan that is essentially correct. For example, in planning to drive to work, one may pick out the route, and modify this due to other complications (say, stopping at a store along the way). all at the level of quite abstract actions such as taking a certain road to a certain place. The actual actions that are executed are not determined until one actually is doing the driving. At that time, one drives and reacts to the other cars on the road, and other unforeseen obstacles that might arise. One could not plan out all these details in advance as they involve information that is unknown until the time of execution. While this may be an extreme example, any realistic problem one examines can be seen to have similar properties.

A probabilistic (or preference-based) model is essential because of the uncertainty mentioned above. Since the planner cannot anticipate all the details

of the actual execution, and something can always occur that makes the plan non-executable (an ice storm might put live power lines across the road, to pick a random example!), one cannot ever prove absolutely that a plan will work. Rather, we can only show that a plan is quite likely to succeed, or that it is the best of the options we have. These notions involve evaluating the relative worth of plans and their likelihood of successful execution.

So resource-bounded planning makes us abandon the simplifying assumptions of having a completely knowable, deterministic world which underlies the majority of planning formalisms so far. In order to be able to formulate plans given more or less time, it is important that solutions can be created incrementally such that at all times, some plan to achieve the goals is present. It may not be a reliable plan, but it is a plan, and if there is no more time, it is the best plan that can be come up with. This is in contrast to many generative planning algorithms that incrementally work through time from the goal to the initial state, or vice versa. For instance, a standard regression-style planner may work backwards from the goal. In this case, the first action to be performed is typically the last action planned! So the entire plan must be built before a single step can be taken. Working forwards in time is no better for one has no idea that the first actions planned will be useful until the final action is planned showing that the sequence achieves the goal.

If the planning algorithm cannot work by incrementally updating time, it must work with actions that span the time between the initial state and the goal. Actions, including abstract actions, must be associated with the following information at the minimum:

- a description of the transitions that the action involves - this information is crucial for determining whether an action is relevant to a specific problem;

- a specification of how the action can be performed, possibly in terms of other abstract actions. This specification should be fairly specific as to the first steps to be taken, but might be quite vague as to how the action is completed.

I view our work on formulating planning as temporal reasoning as developing

a basis for such a representation. The "precondition-effect" representation of action is far too weak to describe most actions reasonably, especially actions that are suitably abstract. Rather, an action description may include conditions that must hold while the action is being executed, may involve partial temporal ordering, and many other complexities. Similarly, the specification of how an action may be performed may be considerably more complex than a sequence of substeps. Some parts of the specification may name specific (abstract) actions, others might just identify other transitions that must be accomplished somehow, the details of which must be worked - either at planning time if there is sufficient time, or at execution time.

Planning = Decision Theory + Execution Monitoring

Steve Hanks

Department of Computer Science & Engineering, FR-35
University of Washington
Seattle, WA 98195

By "resource-bounded planning" I assume what is meant is "planning under time pressure." In some sense all interesting planning involves deciding how to use scarce resources efficiently—otherwise how can we place a value on, or even define the notion of, finding an efficient plan? It seems clear from the workshop focus statement that we are to consider the problems that arise when an agent has to plan at the same time it is executing a plan—the assumption is that both consume the resource of time and they are competing for that resource, either because the agent can't do both at the same time or because postponing commitment to act may force the agent to forego a promising course of action.

I see two sets of issues, one involving representation and one control. The representation problem arises in trying to express the tradeoff in the first place. One must have some notion of the consequences of making these plan/act decisions: what kinds of gains do we expect from thinking, what kinds of losses do we expect if we fail to act. Probability and utility models seem right for this task. Previous and current work of mine involves trying to build these models: how to maintain a probabilistic model of the world and how to build a utility model appropriate for a particular set of goals. Decision theory provides a formally satisfying model for making planning decisions.

Adopting decision theory as a representation has led many to apply it to mediating the plan/act decision as well—in other words, using decision theory to solve the control problem as well as the representation problem. I think this is an inappropriate use of the formalism for two reasons: it's unrealistic to expect that we can do the analysis right, even in principle and *certainly* not in practice, and it leads (or has led) to the fundamentally misguided assumption that planning and acting are mutually exclusive.

My line of research starts with the premise that deliberation and action are *always* going on, and always simultaneously. The problem therefore is one of coordinating the two behaviors rather than deciding which to do at any given moment—the two processes exchange information rather than compete for a processor, as it were. Under this scheme an agent must have the ability to do three things:

1. to commit quickly to the current best course of action when necessary
2. to adapt or improve a plan incrementally if time permits
3. to know when time is indeed critical.

The parts of my system that manage probabilistic reasoning already have the first two properties: they can provide a quick estimate of a fact's probability and improve that estimate given more time. I'm now working on using this system to build a decision-theoretic planner that has the same property.

The third item has to do with how the deliberative agent keeps itself apprised of the changing state of the world. It has to know when it's time to commit, when it has time to think, and when the world has changed enough so it should change strategies or rethink its commitments. There are basically two problems: how does the deliberator get information about the world, and how does it recognize the impact that information has on its plans. This problem is closely related to the one mentioned above of coordinating the behaviors of an autonomous planning system with an autonomous execution system. The deliberator has to know what's being done, and the executor needs to know (roughly) what to do next. I'm currently working with Jim Firby at the University of Chicago on how one passes information back and forth between two such processes, and how knowledge of the evolving world should affect the planning process.

My position, therefore, is that the ability to plan well under time limitations is an emergent property of the system more than a property that can be designed in or a property that can be reasoned about formally and without reference to an implementation.¹ It emerges from the ability to recognize tradeoffs and pick the best option at the moment, and the ability to monitor the environment with an eye to new information, opportunities, and dangers.

¹I don't mean to say that one can't think about the problem during design or talk about the problem in the abstract, rather that we will realize good time-sensitive behavior more easily by observing it in an implementation than we will by proving its existence in a design or formal system.

Searching for Relevance

Josh Tenenber

February 28, 1991

It is well known that expressive planning representations, such as those that are logic based, are expensive to compute with. Typically, only weak search methods that can be brought to bear on finding solutions to problems. Even with less expressive representations, such as table-driven, reactive systems, there can be considerable expense associated with building tables that are both robust, and embody rational decision strategies. In both of these cases, increased knowledge about the problem domain, rather than improving problem solving efficiency, usually results in performance degradation. The problem is that these systems are unable to distinguish between relevant and irrelevant knowledge for a given task. An important theme that emerges, then, is how we might automate the agent's search for relevance in order to improve its performance.

I will discuss research I am involved in concerning statistically based approaches to both the qualification problem in action reasoning and to approximating optimal decision policies in reinforcement learning, how these are affected by the relevance problem, and possible approaches to this problem. Despite the differences in both representation and in the manifestation of computational bottlenecks between the action reasoning and reinforcement learning, there appear to be surprising similarities suggesting that fruitful approaches to one will benefit the other.

Resource-bounded Planning

Marc Vilain

March 10, 1991

Among various approaches to resource-bounded planning, some of the most promising – and most provocative – are based on notions of machine learning. These approaches strive to improve a planner's performance over time by deriving control knowledge from past attempts at solving problems. By analyzing what worked and didn't work in solving a particular problem, these techniques can effectively learn shortcuts towards solving future related problems.

In this talk, I will be covering a number of MITRE projects that have focussed on learning in planning:

- (1) The ULS planning system, which learns probabilistic search control rules through a combination of explanation-based learning and statistical analysis techniques.
- (2) The SMARTPLAN case-based planner, in which learning techniques are currently being used to learn plan modification strategies.
- (3) The ARC constraint-based planner, which uses an inductive classification scheme to learn rules for focussing search in a constraint propagation network.

Piero P. Bonissone and Peter C. Halverson
Artificial Intelligence Laboratory
GE Corporate Research and Development Center
Schenectady, NY

Our work on resource-bounded reasoning processes has concentrated on deductive reasoning systems, identifying issues associated with making inferences in the presence of uncertain and incomplete input data, unpredictable event sequences, and critical time and resource constraints. The reasoning systems used in our architecture have been specifically designed to deal with time-constrained situations described by imprecise, uncertain information:

- The RUM rule-based reasoning system (used in the early design and development stages of the application) provides a representation of uncertain information, uncertainty calculi for inferencing, and selection of calculi for inference control.
- The RUMrunner deployment system extends RUM's mechanisms for reasoning with uncertainty with additional performance and deployment features, eliminating development features which are unnecessary at run-time. Deductive inferencing in RUMrunner is managed by a meta-controller which controls the inference engines to meet real-time performance criteria. The meta-controller consists of an agenda-based task scheduler, to deal with asynchronous input data and queries, and a reasoning planner, which handles reasoning under user-specified time constraints. Reasoning planning is performed by partitioning the deductive inference network into independently-contributing subnets, and selecting some sequence of these subnets for evaluation by the inferencing processes. We have reported on several different schemes for creating and selecting subnets and their effects on complexity, completeness, and performance.
- Using the PRIMO reasoning system (the successor to RUM/RUMrunner), we have analyzed methods for partitioning and scheduling rule networks for execution on parallel architectures.

In addition to continuing our work with these rule-based reasoning systems, we are also involved in other relevant areas, including case-based planning and scheduling using uncertainty, integration of case-based and rule-based planning and reasoning techniques, and real-time temporal event recognition and analysis.

A Retrospective View on the Resource-Bounded Planning Workshop

Steve Hanks

Department of Computer Science & Engineering
University of Washington

I was struck by two aspects of the workshop's discussions: (1) there was pretty much complete agreement as to what are the fundamental issues facing the planning researcher of tomorrow, and (2) there was pretty much no consensus as to what were the best ways (approaches, techniques) to confront those issues.

We agreed on two things: planners must be more expressive, and planners must run faster. The former means that the planner must be able to reason about a larger class of worlds than STRIPS, say, could handle. Richer notions of time, action, change, preference, and so on, are called for. The latter is a reaction to the horrible computational complexity (both in theory and in practice) of current planners. The pity (or did I mean to say "the challenge") is that these two requirements are rather directly at odds.

We were thus united in the view of what the ideal next-generation planner should do, and even on the sorts of domains in which the next-generation planner should operate. The "Trainworld" domain Allen and his group are working on bears a striking similarity to the "Truckworld" domain I've been working on¹, and both make the same sorts of simplifying assumptions that might be brought to the study of planning in realistic transportation or manufacturing domains.

We saw a striking lack of consensus in our opinions on how we are to realize this planner—what issues should be pursued next, and what techniques are likely to lead to progress. Perhaps there was more agreement in the area of how to make the planner more expressive: the answer seems to be decision theory plus temporal logic. I don't consider this an answer, at least not in and of itself: the decision-theoretic formalism is not a panacea any more than first-order logic was a panacea—it represents a language suitable for stating the difficult problems, but is not itself a solution to those problems. Indeed, as I noted in my position paper, there is considerable disagreement as to where and how the technique should be applied. The representation problems have just begun to be explored, and the computational problems loom large.

Two other proposals were presented for building faster planners: abandoning the traditional method of generating plans from first principles, and improving performance by letting the planner learn from experience. Vilain presented some work on case-based planning, which in opposition to rule-based planning, takes the view that new situations will tend to look mostly like old situations,

¹Though the two were conceived of and developed separately

thus it will be advantageous to re-use plans that worked in those old situations. Tenenbergs talked about two sorts of learning: reinforcement learning, whereby an agent learns more effective strategies by refining its old ones in response to feedback, and also statistical learning about the success of actions.

One theme that I tried to champion was that timely behavior is ultimately realized (or not) at execution time, so the plan's execution system must play a vital role in realizing that timely behavior. This position is more in line with the "reactive planning" point of view, of which we heard surprisingly little. I argue that the device that is executing the plan can play two valuable roles: it can cope with low-level uncertainties that will inevitably appear at execution time, thus allowing the planner to reason with a higher-level view of the world (and thus plan more quickly). Second, it can provide the planner with more information about the world as it executes.

Obviously all of these techniques show promise, and, just as obviously, none of them in and of themselves constitute a solution. The case-based view that new situations resemble old, for example, is a powerful technique for facilitating planning at a high level, but at lower levels this assumption breaks down completely. (No new situation will resemble an old one if you look at it in fine enough detail.) Reinforcement learning, on the other hand, is a powerful way to generate low-level, single-purpose behaviors, but probably cannot be used to train an innovative high-level planner.

Our success in building an intelligent agent depends crucially on our ability to combine these techniques fruitfully. Perhaps this is an obvious insight, but it is one that has as been almost completely ignored in the actual development of systems. (Consider, for example, the work on reactive planning, which has been conducted, for the most part, without attention to how the so-called reactive agent might be imbedded in a larger context.) As we set our sights higher, and attempt to capture a greater range of intelligent behavior, we can less and less afford to ignore it. We have made much more progress on the the pieces than we have on the whole, and it's time to reverse that process.

Participants

James Allen	james@cs.rochester.edu
Piero Bonissone	bonisson@sol.crd.ge.com
Tom Clark	Clarkt@lonex.radc.af.mil
John Crowter	crowter@aovax.radc.af.mil
Jeff Davis	davis@aivax.radc.af.mil
Doug Dyer	dyer@aivax.radc.af.mil
Northrup Fowler	fowler@aivax.radc.af.mil
Al Frantz	frantz@aivax.radc.af.mil
Peter Halverson	halverson@sol.crd.ge.com
Steve Hanks	hanks@cs.washington.edu
Leo Hartman	leo@cs.rochester.edu
Lou Hoebel	hoebel@aivax.radc.af.mil
Robert Kruchten	Kruchtenr@lonex.radc.af.mil
Ray Liuzzi	Liuzzir@lonex.radc.af.mil
Nat Martin	martin@cs.rochester.edu
Clark Mollenhauer	
Massimo Poesio	poesio@cs.rochester.edu
Nancy Roberts	nancyr@aivax.radc.af.mil
Robert Ruberti	ruberti@aivax.radc.af.mil
Karen Sidoren	karens@aivax.radc.af.mil
Jennifer Skidmore	skiddles@ra.radc.af.mil
Frank Sliwa	Sliwa@aivax.radc.af.mil
Josh Tenenberg	josh@cs.rochester.edu
David Traum	traum@cs.rochester.edu
Marc Vilain	mbv@linus.mitre.org
Roger Weber	

References
The March 1991 Workshop on Resource-Bounded Planning
Rome Laboratory
Rome NY

James F. Allen

Planning as Temporal Reasoning

Proc. of the Second International Conference on Principles of Knowledge Representation and Reasoning (1991)

James F. Allen, Johannes A. Koomen

Planning Using a Temporal World Model

IJCAI 83

Piero P. Bonissone, Peter C. Halverson

Time-Constrained Reasoning Under Uncertainty

Journal of Real-Time Systems, 2, pp. 25-45 (1990)

Steve Hanks

Building Intelligent Agents

Working paper, Dept. of Computer Science and Engineering, University of Washington, Seattle WA (1991).

Steve Hanks

Controlling Inference in Planning Systems: Who, What, When, Why and How

Technical Report #90-04-01, Dept. of Computer Science and Engineering, University of Washington, Seattle WA (1991)

Josh D. Tenenbergh, Jay C. Weber

A Statistical Approach to the Qualification Problem

Working paper, Computer Science, University of Rochester

Leo B. Hartman, Josh D. Tenenbergh

Performance in Practical Problem Solving

IJCAI 87

Dennis Connolly, Steve Coley, Phyllis Koton, Stuart McAlpin,

Alice Mulvehill, Marc Vilain

Learning Representation by Integrating Case-based and Inductive Learning

Submitted to the Machine Learning Workshop 1991

David S. Day

Learning Variable Descriptors for Applying Heuristics Across CSP Problems

Submitted to the Machine Learning Workshop 1991

Melissa P. Chase, Monte Zweben, Richard L. Piazza, John D. Burger,

Paul P. Maglio, Haym Hirsh

Approximating Learned Search Control Knowledge

Machine Learning Workshop -1989

**MISSION
OF
ROME LABORATORY**

Rome Laboratory plans and executes an interdisciplinary program in research, development, test, and technology transition in support of Air Force Command, Control, Communications and Intelligence (C³I) activities for all Air Force platforms. It also executes selected acquisition programs in several areas of expertise. Technical and engineering support within areas of competence is provided to ESD Program Offices (POs) and other ESD elements to perform effective acquisition of C³I systems. In addition, Rome Laboratory's technology supports other AFSC Product Divisions, the Air Force user community, and other DOD and non-DOD agencies. Rome Laboratory maintains technical competence and research programs in areas including, but not limited to, communications, command and control, battle management, intelligence information processing, computational sciences and software producibility, wide area surveillance/sensors, signal processing, solid state sciences, photonics, electromagnetic technology, superconductivity, and electronic reliability/maintainability and testability.