

AFRL-IF-RS-TR-2004-313
Final Technical Report
November 2004



LAYERED MULTI-TEMPLATE RETRIEVAL, ADAPTATION AND LEARNING

Carnegie Mellon University

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. K443

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

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

AFRL-IF-RS-TR-2004-313 has been reviewed and is approved for publication.

APPROVED: /s/

DALE W. RICHARDS
Project Engineer

FOR THE DIRECTOR: /s/

JAMES A. COLLINS, Acting Chief
Information Technology Division
Information Directorate

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 074-0188	
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 this 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 NOVEMBER 2004		3. REPORT TYPE AND DATES COVERED Final Jul 00 –Sep 03
4. TITLE AND SUBTITLE LAYERED MULTI-TEMPLATE RETRIEVAL, ADAPTATION AND LEARNING			5. FUNDING NUMBERS C - F30602-00-2-0549 PE - 63760E PR - ATEM TA - P0 WU - 11	
6. AUTHOR(S) Manuela Veloso				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University 5000 Forbes Avenue Pittsburgh Pennsylvania 15213-3890			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency AFRL/IFT 3701 North Fairfax Drive Arlington Virginia 22203-1714			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2004-313	
11. SUPPLEMENTARY NOTES AFRL Project Engineer: Dale W. Richards/IFT/(315) 330-3014/ Dale.Richards@rl.af.mil				
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.				12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 Words) This effort was part of the DARPA Active Templates program (2000-2004) to revolutionize mission planning, mission execution, and related command and control processes. Extensive use is made of previous research in generative planning and learning, case-based and mixed-initiative plan adaptation, real-time integration of action and execution, and multi-agent control and learning. Technology was developed to support users in creating and managing template-based plan, allowing them to anticipate multiple contingencies and dynamically re-plan based on real-time sensory information. The research is grouped into the following themes: (1) Allocation of communications spectrum frequencies, (2) Extraction of plan rationale and the learning of planning templates, (3) New abstraction techniques for reinforcement learning to improve the efficiency of automatic control algorithms, (4) Opponent modeling in dynamic multi-agent environment, (5) Multi-agent learning and limitations, and (6) Planning using symbolic model-based techniques. An extensive bibliography is included listing publications which describe the results of these research tasks in more detail.				
14. SUBJECT TERMS Automated Planning, Reinforcement Learning, Intelligent Agents, Templates			15. NUMBER OF PAGES 20	
			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 UL	

Table of Contents

Overview of the Active Templates Program	1
Layered, Multi-Template Retrieval, Adaptation and Learning	2
Technical Challenges	2
Technical Approach	3
Communication Frequency Planner	3
Learning Planning Templates: Domain-Specific Planners	4
Use of Macros in Control Learning for Complex Domains	5
Opponent Modeling and Coaching	6
Multi-Agent Control Learning	7
Efficient Planning using Symbolic Model-Based Techniques	8
Publications	10

Overview of the Active Templates Program

The goal of the DARPA Active Templates program (2000-2004) has been to develop technologies designed to revolutionize aspects of mission planning, mission execution, and related command and control processes in order to improve the ability of military units to organize, plan and conduct military operations. A particular emphasis has been placed on demonstrating the application of these technologies to problems faced in the course of conducting special operations and other time-critical, small-unit missions. These prototype technologies have focused on the temporal and spatial aspects of the associated mission information, as well as on the supporting infrastructure necessary to implement them in a distributive, collaborative environment. Ordinary users must be able to dynamically define their own workflow systems, interface to disparate data sources, and specify problems to be solved by advanced problem-solvers.

Technologies like electronic mail and the web have significant advantages and have had huge impact on military capability, but there is technology beyond the e-mail and the web that promises even more. The objective of this program was to create a new structured way of communicating and operating that enables the computer to build a partial model of the situation. With this model, the computer will be able to assist users in:

- finding and integrating information
- prioritizing requests and tasks
- determining key information elements and available options
- reuse similar past solutions or defaults
- coordinating decisions and routing information to others
- recognizing important changes

This new technology will provide a means, using tailored templates, for doing common tasks and sharing data in a dynamic workflow system. These templates will provide prioritized information that updates in real time and triggers the computer to automatically analyze the impact of changes, suggest default actions, auto-coordinate decisions, and capture a digital history for purposes of accountability, training, and process improvement. Most importantly,

when we sketch out plans (for example), we'll have share-able data, and this data is necessary to move beyond the spell-checking, key-word search level of automation in place today.

Layered, Multi-Template Retrieval, Adaptation and Learning

Work on this effort was based extensively on previous research in generative planning and learning, case-based and mixed-initiative plan adaptation, real-time integration of action and execution, and multi-agent control and learning. Users performing complex planning tasks should be able to rely upon fully-automated agent-based software tools – while being able to fully inspect, guide, and alter the software agents' actions. Critical to this is the need for a framework for creating and managing template plans that can anticipate multiple contingencies and dynamically re-plan based on real-time sensory information, both based on the system's own execution and on unanticipated adversarial actions. Templates must be able to improve incrementally and new templates generated based on past template-filling episodes.

Technical Challenges

Simplicity of use: Users must be able to examine the Special Operations Forces (SOF) adaptation of templates in a layered approach corresponding to increasingly higher levels of detail.

Incremental and dynamic template adaptation: The “template-filling” and adaptation process must be approached in a mixed-initiative, incremental, and rationale-based way in response to dynamically perceived relevant world changes.

Multi-template management: Retrieval, adaptation, and introspection of multiple templates are necessary to allow for user-guided identification of resource contentions and other conflicts. Based on template-filling behaviors, automated suggestions can be provided to assist the user in merging multiple templates.

Template generation and learning: Based on the user's input and changes during dynamic execution of a plan, the system should be able to learn templates of increasing quality. Abstract templates and partially instantiated templates can enable the generation of plan structures

involving multiple taskforces and multiple decision points, the data structures for monitoring the execution of those plans.

Technical Approach

The approach taken by CMU to address the underlying technical challenges associated with making Active Templates a viable Command and Control concept eventually came to focus on six themes: (1) Allocation of communications spectrum frequencies, (2) Extraction of plan rational and the learning of planning templates, (3) New abstraction techniques for reinforcement learning to improve the efficiency of automatic control algorithms, (4) Opponent modeling in dynamic multi-agent environment, (5) Multi-agent learning and limitations, and (6) Planning using symbolic model-based techniques. Each of these areas is described in the following section of this report. An extensive bibliography at the end of this report lists publications which describe the results of these research tasks in more detail.

Communication Frequency Planner

A frequency resource allocation system, *CommPlanner*, was developed for use in the Special Forces domain. The first *CommPlanner* prototype calculated what frequency bands are available in an area of interest by generating a frequency channel usage list in the area of interest. The channel usage list is generated by parsing a given communication resources database for known transmitters and receivers in the area of interest. Once the channel usage list has been generated, *CommPlanner* finds and allocates free frequency bands meeting the specified requirements of bandwidth, frequency range and tolerance.

The *CommPlanner* prototype operates in a stand-alone manner with a graphical interface and a nominal database. Operation of the *CommPlanner* prototype was successfully demonstrated to the DARPA Program Manager, SOF DARPA consultants, and to operational users who currently perform this tedious task by hand. Interaction with all the above parties led to a number of revisions to the prototype.

The revised *CommPlanner* program is a complete piece of software that can be integrated with other SOF tools, most notably the *C2PC* program. *CommPlanner* takes as input specification files that define how many frequency bands are required, their frequency range, and in what rectangular area (as defined by two pair of longitude and latitude coordinates). Using information contained in the XML input file, *CommPlanner* then queries the database to retrieve the relevant transmitter and receiver information that is within the specified frequency range and area of interest. The next step of the process, which is currently under development, is to generate the frequency assignments using the previous code as a starting point and to convert it to an appropriate XML output format. *CommPlanner* can also communicate with real database files, e.g., *FCCReg1.mdb*. *CommPlanner* can retrieve the data record set from the database with the information originated from the XML file. So the basic information flow is for *CommPlanner* to fetch the specification from the input XML file, retrieve the record set from the database with the fetched specification, do the frequency band calculation based on the record set and produces an output file in the legitimate XML file format. *CommPlanner* is self-contained, with a simple GUI by which the end user may interact.

Learning Planning Templates: Domain-Specific Planners

One of the key issues identified in template learning from examples is the ability to “understand” the reasons or rationale for the steps in a given execution example.

Several tasks, such as plan reuse and agent modeling, need to interpret a given or observed valid plan to generate the underlying plan rationale. Although there have been several successful contributions to this rationale-extraction problem, they do not apply to complex plans, in particular to plans with actions that have conditional effects. *SPRAWL*, an algorithm to find a minimally-annotated, partially-ordered structure in an observed totally-ordered plan with conditional effects was developed. The algorithm proceeds in a two-phased approach. First the given plan is preprocessed using a novel “needs analysis” technique that builds a “needs tree” to identify the dependencies that link all the literals, including the conditional effects, in the totally ordered plan. This needs tree is then further processed to construct a partial ordering that captures the complete rationale of the given plan.

This capability was then enhanced to support automatically acquiring templates from example plans. As is well known, general-purpose planners can solve problems in a variety of domains but can be quite inefficient in a given domain. Domain-specific planners, on the contrary, are more efficient, but are difficult to create. Template-based planning has been introduced as a novel paradigm for automatically generating domain-specific programs, or templates. *DISTILL*, an algorithm for learning templates automatically from example plans, was thus developed. *DISTILL* converts a given example plan into a template and then merges it with previously learned templates. Empirical results in simple domains have been achieved that show that the templates automatically learned by *DISTILL* compactly represent its domain-specific planning experience. Furthermore, the templates situationally generalize the given example plan, thus allowing them to efficiently solve problems that were not previously encountered.

This research is core to the Ph.D. Thesis of Elly Winner, expected to be finished in the fall 2005, entitled “*Learning Domain-Specific Planners from Example Plans.*”

Use of Macros in Control Learning for Complex Domains

Methods for speeding up automatic control algorithms have been investigated, specifically, new abstraction techniques for Reinforcement Learning and Semi-Markov Decision Processes (SMDPs). The use of policies as temporally abstract actions is introduced. This is different from previous definitions of temporally abstract actions as termination criteria are not involved. An approach for processing previously solved problems to extract these policies has been developed utilizing a method for using supplied or extracted policies to guide and speed up problem solving of new problems. Extracting policies are treated as a supervised learning task and a new algorithm, *LUMBERJACK*, that extracts repeated sub-structure within a decision tree was defined. Another algorithm, *TTREE*, that combines state and temporal abstraction to increase problem solving speed on new problems was also introduced. *TTREE* solves SMDPs by using both user and machine supplied policies as temporally abstract actions while generating its own tree-based abstract state representation. By combining state and temporal abstraction in this way, *TTREE* is the only known SMDP algorithm that is able to ignore irrelevant or harmful sub-regions within a supplied abstract action while still making use of other parts of the abstract action.

This research led to the Ph.D. Thesis of William Uther, finished in August 2002, entitled “*Tree-Based Hierarchical Reinforcement Learning*.”

Opponent Modeling and Coaching

A set of algorithms have been developed which together are able to process observation logs of continuous, dynamic multi-agent movements and actions in complex environments. The algorithms analyze these logs and extract sequential templates: i.e., important repeated instances of sequential behavior by the observed agents. In other words, given a set of recordings of the movements and actions of multiple agents the developed algorithms can extract repeating patterns of sequential interaction. These algorithms have been fully implemented in the context of the RoboCup Soccer domain.

The algorithms work in two stages: The first stage transforms a stream of continuous world-state observations, including object and agent positions, into a stream of categorical observations corresponding to recognized atomic single-agent behaviors. These algorithms rely on domain knowledge to identify instances of recognizable behavior, matching the continuous observation stream against conditions that check for agent movement over time, relative positions of agents and objects, and spatial patterns.

Once this stream of recognized behaviors is available, a second algorithm is used to identify repeating patterns. First, the recognized behavior stream is segmented by team ID so that the behaviors of different teams are separated. The segments of each team are inserted into a separate *trie* structure, a tree-like data structure that supports efficient storage and access, while maintaining count of how many times each sequence and sub-sequence appeared. Finally, the *trie* is traversed to assign a rank to each sequence by means of a ranking function. These ranks can then be sorted for the final output.

The problem of modeling observed execution has also been extensively investigated. An advising agent, a coach, provides advice to other agents about how to act. This is accomplished by an advice generation method using observations of agents acting in an environment. Given an abstract state definition and partially specified abstract actions, the algorithm extracts a Markov Chain, infers a Markov Decision Process (MDP), and then solves the MDP (given an arbitrary

reward signal) to generate advice. This capability has been evaluated in a simulated robot-soccer environment and experimental results show improved agent performance when using the advice generated from the MDP for both a sub-task and the full soccer game.

This research is core to the Ph.D. Thesis of Patrick Riley, expected to be finished in the Fall 2005, entitled “*Advice Generation from Modeling Observed Execution.*”

Multi-Agent Control Learning

Learning to act in a multi-agent environment is a challenging problem. Optimal behavior for one agent depends upon the behavior of the other agents, which are learning as well. Multi-agent environments are therefore non-stationary - violating the traditional assumption underlying single-agent learning. In addition, agents in complex tasks may have limitations, such as physical constraints or designer-imposed approximations of the task that make learning tractable. Limitations prevent agents from acting optimally, which complicates the already challenging problem. A learning agent must effectively compensate for its own limitations while exploiting the limitations of the other agents. This research theme focuses on these two challenges, namely multi-agent learning and limitations, and includes four main contributions.

First, the thesis introduces the novel concepts of a variable learning rate and the *WoLF* (Win or Learn Fast) principle to account for other learning agents. The *WoLF* principle is capable of making rational learning algorithms converge to optimal policies. By doing so it achieves two properties, rationality and convergence, which have not been achieved by previous techniques. The converging effect of *WoLF* has been proven for a class of matrix games, and demonstrated empirically for a wide-range of stochastic games.

Second, the thesis contributes an analysis of the effect of limitations on the game-theoretic concept of Nash Equilibria. The existence of equilibria is important if multi-agent learning techniques, which often depend on the concept, are to be applied to realistic problems where limitations are unavoidable. The thesis introduces a general model for the effect of limitations on agent behavior, which is used to analyze the resulting impact on equilibria. The thesis shows that equilibria do exist for a few restricted classes of games and limitations, but even well-behaved limitations do not preserve the existence of equilibria in general.

Third, the thesis introduces *GraWoLF*, a general-purpose, scalable, multi-agent learning algorithm. *GraWoLF* combines policy gradient learning techniques with the *WoLF* variable learning rate. The effectiveness of the learning algorithm has been demonstrated in both a card game with an intractably large state space, and an adversarial robot task. These two tasks are complex and agent limitations are prevalent in both.

Fourth, the thesis describes the *CMDragons* robot soccer team strategy for adapting to an unknown opponent. The strategy uses a notion of plays as coordinated team plans. The selection of team plans is the decision point for adapting the team to its current opponent, based on the outcome of previously executed plays. The *CMDragons* were the first RoboCup robot soccer team to employ online learning to autonomously alter its behavior during the course of a game.

These four contributions demonstrate that it is possible to effectively learn to act in the presence of other learning agents in complex domains when agents may have limitations. The introduced learning techniques are proven effective in a class of small games, and demonstrated empirically across a wide range of settings that increase in complexity.

This research led to the Ph.D. Thesis of Michael Bowling, finished in August 2003, entitled “*Multi-Agent Learning in the Presence of Agents with Limitations.*”

Efficient Planning using Symbolic Model-Based Techniques

Automated planning considers selecting and sequencing actions in order to change the state of a discrete system from some initial state to some goal state. This problem is fundamental in a wide range of industrial and academic fields including robotics, automation, embedded systems, and operational research. Planning with non-deterministic actions can be used to model dynamic environments and alternative action behavior. One of the currently best known approaches is to employ reduced ordered Binary Decision Diagrams (BDDs) to represent and generate plans using techniques developed in symbolic model checking. However, the approach is challenged by a frequent blow-up of the BDDs representing the search frontier and a limited number of solution classes.

This thesis addresses both of these problems. With respect to the first, it contributes a general

framework called *state-set branching* that seamlessly combines classical heuristic search and BDD-based search. Our experimental results show that the performance of state-set branching often dominates both blind BDD-based search and ordinary heuristic search. In addition, it consistently outperforms any previous approach to guide a BDD-based search of which we are aware. We show that state-set branching naturally generalizes to non-deterministic planning and introduce heuristically guided versions of the current BDD-based non-deterministic planning algorithms.

With respect to the second problem, the thesis introduces two frameworks called *fault tolerant planning* and *adversarial planning*. Fault tolerant planning addresses domains where non-determinism is caused by rare errors. The current solution classes handle this situation poorly by taking all fault combinations into account or produce solutions which are too weak. The thesis contributes a new class of solutions called *fault tolerant plans* that are robust to a limited number of faults. In addition, it introduces specialized BDD-based algorithms for synthesizing fault tolerant plans.

Adversarial planning considers situations where non-determinism is caused by uncontrollable, but known, environment actions. The current solution classes of BDD-based non-deterministic planning assume a “friendly” environment and may never reach a goal state if the environment is hostile and informed. The thesis contributes efficient BDD-based algorithms for synthesizing winning strategies for such problems.

This research led to the Ph.D. Thesis of Rune Jensen, finished in August 2003, entitled “*Efficient BDD-Based Planning for Non-Deterministic, Fault-Tolerant, and Adversarial Domains.*”

Publications

1. Bowling, Michael. “Convergence Problems of General Sum Multi-Agent Reinforcement Learning.” Proceedings of the Seventeenth International Conference on Machine Learning (CML-2002), June 2000.
2. Bowling, Michael, and Manuela Veloso, An Analysis of Stochastic Game Theory for Multi-Agent Reinforcement Learning. Carnegie Mellon University Technical Report, CMU-CS-00-165, 2000.
3. Bowling, Michael, and Manuela Veloso. “Convergence of Gradient Dynamics with a Variable Learning Rate.” Proceedings of International Conference on Machine Learning (CML-2001), pp. 27–34, Williams College, MA, June 2001.
<http://www.cs.cmu.edu/~mmv/papers/01icml-mike.pdf>
4. Bowling, Michael, and Manuela Veloso. “Rational Learning of Mixed Equilibria in Stochastic Games.” Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-2001), pp. 1021–1026, August 2001. <http://www.cs.cmu.edu/~mmv/papers/01ijcai-mike.pdf>
5. Bowling, Michael, and Manuela Veloso. “Rational and Convergent Multi-Agent Learning.” Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-2001), August 2001.
6. Bowling, Michael and Manuela Veloso. “Multi-Agent Learning Using a Variable Learning Rate.” International Conference on Machine Learning (ICML), 2001.
7. Browning, Brett, Gal Kaminka, and Manuela Veloso. “Principled Monitoring of Distributed Agents for Detection of Coordination Failure.” Proceedings of the Seventh International Symposium on Distributed Autonomous Robotic Systems (DARS), June 2002.

8. Bowling, Michael and Manuela Veloso. “Scalable Learning in Stochastic Games.” Proceedings of the National Conference on Artificial Intelligence (AAAI-02) Workshop on Game Theoretic and Decision Theoretic Agents, Edmonton, Canada, August 2002.
<http://www.cs.cmu.edu/~mmv/papers/02aaaiws-gtdt.pdf>
9. Bowling, Michael and Manuela Veloso. “Multi-Agent Learning Using a Variable Learning Rate.” Artificial Intelligence, 136:215–250, 2002.
<http://www.cs.cmu.edu/~mmv/papers/02aij-mike.pdf>
10. Bowling, Michael and Manuela Veloso. “Simultaneous Adversarial Multi-Robot Learning.” Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI’03), Acapulco, Mexico, August 2003. <http://www.cs.cmu.edu/~mmv/papers/03ijcai-grawolf.pdf>
11. Bowling, Michael, Rune Jensen, and Manuela Veloso. “A Formalization of Equilibria for Multi-Agent Planning.” Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI’03), Acapulco, Mexico, August 2003.
<http://www.cs.cmu.edu/~mmv/papers/03ijcai-map.pdf>
12. Bowling, Michael. Multi-Agent Learning In the Presence of Agents with Limitations. Carnegie Mellon University PhD Dissertation, CMU-CS-03-118, 2003.
13. Bowling, Michael, Brett Browning, and Manuela Veloso. “Plays as Team Plans for Coordination and Adaptation.” Proceedings of the 14th International Conference on Automated Planning and Scheduling (ICAPS-04), Vancouver, June 2004.
<http://www.cs.cmu.edu/~mmv/papers/04icaps-plays.pdf>
14. Hiyakumoto, Laurie, and Manuela Veloso. “Planning and Execution for Information Retrieval.” Proceedings of the European Conference on Planning (ECP ’01), 2001.
15. Hiyakumoto, Laurie, and Manuela Veloso. “Towards Planning and Execution for Information Retrieval.” Proceedings of the Sixth International Conference on Artificial Intelligence Planning Systems (AIPS ’02), Workshop on Exploring Real-World Planning, Toulouse, April 2002. <http://www.cs.cmu.edu/~mmv/papers/02aipsw-laurie.pdf>

16. Jensen, Rune M. and Manuela Veloso. "OBDD-Based Universal Planning: Specifying and Solving Planning Problems For Synchronized Agents In Non-Deterministic Domains." Artificial Intelligence Today: Recent Trends and Developments. Ed. Wooldrige, M. and Manuela Veloso. pp. 212–248. Springer-Verlag, 1999.
<http://www.cs.cmu.edu/~mmv/papers/UMOP-AItoday.pdf>
17. Jensen, Rune M. and Manuela Veloso. "OBDD-Based Universal Planning For Synchronized Agents In Non-Deterministic Domains." Journal of Artificial Intelligence Research, 13:189– 226, July 2000. <http://www.cs.cmu.edu/~mmv/papers/jair-umop.pdf>
18. Jensen, Rune, Manuela Veloso, and Michael Bowling. "Optimistic and Strong Cyclic Adversarial Planning." Proceedings of the European Conference on Planning (ECP-2001), Toledo, Spain, October 2001. <http://www.cs.cmu.edu/~mmv/papers/01ecp-rune.pdf>
19. Jensen, Rune M., Randy E. Bryant, and Manuela Veloso. "SetA*: An efficient BDD-Based Heuristic Search Algorithm." Proceedings of the National Conference on Artificial Intelligence (AAAI), 2002.
20. Jensen, Rune M., Randy E. Bryant, and Manuela Veloso. "An efficient BDD-based A* Algorithm." Proceedings of the Sixth International Conference on Artificial Intelligence Planning and Scheduling (AIPS), Workshop on Planning via Model Checking, 2002.
21. Jensen, Rune M., Randy E. Bryant, and Manuela Veloso. "SetA* Applied to Channel Routing." Carnegie Mellon University Technical Report, CMU-CS-02-172, 2002.
22. Jensen, Rune, Manuela Veloso, and Randy Bryant. "Guided Symbolic Universal Planning." Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS'03), Trento, Italy, June 2003. <http://www.cs.cmu.edu/~mmv/papers/03icaps-rune2.pdf>
23. Jensen, Rune, Manuela Veloso, and Randy Bryant. "Synthesis of Fault-Tolerant Plans for Non-Deterministic Domains." Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS'03) Workshop on Planning under Uncertainty and Incomplete Information, Trento, Italy, June 2003. <http://www.cs.cmu.edu/~mmv/papers/03icaps-rune1.pdf>

24. Jensen, Rune M. Efficient BDD-Based Planning for Non-Deterministic, Fault-Tolerant and Adversarial Domains. Carnegie Mellon University PhD Dissertation, CMU-CS-03-139, 2003.
25. Jensen, Rune M., Randy E. Bryant, and Manuela Veloso. "State-Set Branching: Leveraging OBDDs for Heuristic Search." Artificial Intelligence, 2003, (Under Revision). <http://www.cs.cmu.edu/~mmv/papers/03aij-rune.pdf>
26. Jensen, Rune M., Manuela Veloso and Randy E. Bryant. "Guided Symbolic Universal Planning." Proceedings of the 13th International Conference on Automated Planning and Scheduling (ICAPS-03), 2003.
27. Jensen, Rune M., Manuela Veloso and Randy E. Bryant. "Fault Tolerant Planning: Toward Probabilistic Uncertainty Models in Symbolic Non-Deterministic Planning." Proceedings of the 14th International Conference on Automated Planning and Scheduling (ICAPS-04), Vancouver, June 2004. <http://www.cs.cmu.edu/~mmv/papers/04icaps-rune.pdf>
28. Kaminka, Gal A., "Towards Robust Teams with Many Agents." Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS), October 2000
29. Kaminka, G. A., Manuela Veloso, et al. "Gamebots: The Ever-Challenging Multi-Agent Research Test-Bed." Communications of the ACM, January 2002. <http://www.cs.cmu.edu/~mmv/papers/02gamebots.pdf>
30. Kaminka, Gal, Mehmet Fidanboyly, Allen Chang, and Manuela Veloso. "Learning the Sequential Coordinated Behavior of Teams from Observations." Proceedings of the RoboCup-2002 Symposium, Fukuoka, Japan, June 2002. <http://www.cs.cmu.edu/~mmv/papers/02RoboCup-sequence.pdf>
31. Riley, Patrick, Peter Stone, and Manuela Veloso. "Layered Disclosure: Revealing Agents' Internals." Proceedings of the International Workshop on Agent Theories, Architectures, and Languages (ATAL-2000), Boston, July 2000. <http://www.cs.cmu.edu/~mmv/papers/atal-disclosure.pdf>

32. Riley, Patrick, and Manuela Veloso. "On Behavior Classification in Adversarial Environments." Proceedings of the Fifth International Symposium on Distributed Autonomous Robotic Systems (DARS-2000), October 2000. <http://www.cs.cmu.edu/~mmv/papers/dars00-pat.pdf>
33. Riley, Patrick and Manuela Veloso. "Team Planning for Distributed Execution." Proceedings of the European Conference on Planning (ECP), 2001.
34. Riley, Patrick, and Manuela Veloso. "Planning for Distributed Execution Through Use of Probabilistic Opponent Models." Proceedings of the Sixth International Conference on Artificial Intelligence Planning Systems (AIPS '02), Toulouse, France, April 2002. <http://www.cs.cmu.edu/~mmv/papers/02aips-pat.pdf>
35. Riley, Patrick, Manuela Veloso, and Gal Kaminka. "An Empirical Study of Coaching." Proceedings of the Seventh International Symposium on Distributed Autonomous Robotic Systems (DARS-2002), Fukuoka, Japan, June 2002. <http://www.cs.cmu.edu/~mmv/papers/02dars-coaching.pdf>
36. Riley, Patrick and Manuela Veloso. "Recognizing Probabilistic Opponent Movement Models." RoboCup-2001: The Fifth RoboCup Competitions and Conferences. Ed. A. Birk, S. Coradeschi, and S. Tadokoro. Springer-Verlag, Berlin, 2002. <http://www.cs.cmu.edu/~mmv/papers/LNAI01-coach.pdf>
37. Riley, Patrick, Manuela Veloso, and Gal Kaminka. "Towards Any-Team Coaching in Adversarial Domains." Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'02), pp. 1145–1146, 2002. <http://www.cs.cmu.edu/~mmv/papers/02agents-pat.pdf>
38. Riley, Patrick and Manuela Veloso. "Planning For Distributed Team Execution Using Probabilistic Opponent Models." Journal of Artificial Intelligence Research, 2003 (Under Submission). <http://www.cs.cmu.edu/~mmv/papers/03jair-pat.pdf>

39. Riley, Patrick and Manuela Veloso. "An Overview of Coaching with Limitations." Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'03), pp. 1110–1111, Melbourne, Australia, July 2003.
<http://www.cs.cmu.edu/~mmv/papers/03aamas-pat.pdf>
40. Riley, Patrick and Manuela Veloso. "Advice Generation from Observed Execution: Abstract Markov Decision Process Learning." Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI-04), San Jose, July 2004.
<http://www.cs.cmu.edu/~mmv/papers/04aaai-pat.pdf>
41. Uther, William, and Manuela Veloso. "The Lumberjack Algorithm for Learning Linked Decision Forests." Proceedings of the Symposium on Abstraction, Reformulation and Approximation (SARA 2000), July 2000, Texas.
<http://www.cs.cmu.edu/~mmv/papers/lumberjack-will.pdf>
42. Uther, William, and Manuela Veloso. "The Lumberjack Algorithm for Learning Linked Decision Forests (Condensed)". Proceedings of Pacific Rim International Conferences on Artificial Intelligence (PRICAI-2000), August 2000.
43. Uther, William, and Manuela Veloso. "TTree: Tree-Based State Generalization with Temporally Abstract Actions." Proceedings of the Symposium on Abstraction, Reformulation and Approximation (SARA-2002), Edmonton, Canada, August 2002.
<http://www.cs.cmu.edu/~mmv/papers/02sara-ttree.pdf>
44. Uther, William. Tree-Based Hierarchical Reinforcement Learning, Carnegie Mellon University PhD Dissertation, CMU-CS-02-169, 2002.
45. Uther, William and Manuela Veloso. "TTree: Tree-Based State Generalization with Temporally Abstract Actions." Adaptive Agents and Multi-Agent Systems, LNAI 2636, 266–296. Springer, 2003. <http://www.cs.cmu.edu/~mmv/papers/03will.pdf>

46. Vu, Thuc, Jared Go, Gal Kaminka, Manuela Veloso, and Brett Browning. "MONAD: A Flexible Architecture for Multi-Agent Control." Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'03), Melbourne, Australia, July 2003. <http://www.cs.cmu.edu/~mmv/papers/03aamas-monad.pdf>
47. Winner, Elly, and Manuela Veloso. "Multi-Fidelity Behaviors: Acting with Variable State Information." Proceedings of the National Conference on Artificial Intelligence (AAAI-00), August 2000. <http://www.cs.cmu.edu/~mmv/papers/aaai00-elly.pdf>
48. Winner, Elly and Manuela Veloso. "Plan Structure Extraction." Proceedings of the European Conference on Planning (ECP), 2001.
49. Winner, Elly, and Manuela Veloso. "Automatically Acquiring Planning Templates from Example Plans. Proceedings of the Sixth International Conference on Artificial Intelligence Planning Systems (AIPS '02), Workshop on Exploring Real-World Planning, Toulouse, April 2002. <http://www.cs.cmu.edu/~mmv/papers/02aipsw-elly.pdf>
50. Winner, Elly, and Manuela Veloso. "Analyzing Plans with Conditional Effects." Proceedings of the Sixth International Conference on Artificial Intelligence Planning Systems (AIPS '02), Toulouse, France, April 2002. <http://www.cs.cmu.edu/~mmv/papers/02aips-elly.pdf>
51. Winner, Elly and Manuela Veloso. "DISTILL: Towards Learning Domain-Specific Planners by Example." Proceedings of the International Conference on Machine Learning (ICML'03), Washington, DC, August 2003. <http://www.cs.cmu.edu/~mmv/papers/03icml-elly.pdf>
52. Winner, Elly and Manuela Veloso. "Learning Looping Planners From Examples." The IEEE Intelligent Systems, Special Issue: Agent-Based Intelligent Forms, 2004, (Under Submission). <http://www.cs.cmu.edu/~mmv/papers/04dsPlanners.pdf>