

Curriculum Vita

Prasad Tadepalli

Department of Computer Science

Oregon State University

Corvallis, OR 97331-3902, U.S.A.

Ph: (541)-737-5552 (tadepalli@cs.orst.edu)

Education:

B.Tech (1974-79) in Electrical Engineering, Regional Engineering College, Warangal, India.

M.Tech (1979-81) in Computer Science, I.I.T., Madras, India.

Ph.D (1984-89) in Computer Science, Rutgers University, New Brunswick, NJ, U.S.A.

Positions held:

2009-current Professor, Department of Computer Science, Oregon State University, Corvallis, OR.

1995-2009: Associate Professor, Department of Computer Science, Oregon State University, Corvallis, OR.

1997-98: Visiting Professor, University of California, Berkeley, CA.

1989-95: Assistant Professor, Department of Computer Science, Oregon State University, Corvallis, OR.

1986-89: Research Assistant, Carnegie Mellon University, Pittsburgh, PA.

1984-86: Teaching/Research Assistant, Rutgers University, New Brunswick, NJ.

1983-84: Teaching Assistant, University of Massachusetts, Amherst, MA.

1981-83: Systems Analyst, Bowery Savings Bank, New York, NY.

Ph.D thesis:

“Tractable Learning and Planning in Games,” Rutgers University, New Brunswick, NJ, U.S.A., 1990.

Advisor: Tom Mitchell

Graduate Courses Taught:

Introduction to Artificial Intelligence, Advanced Artificial Intelligence, Machine Learning, Algorithms and Data Structures, Theory of Computation and Formal Languages, Special Topics in Artificial Intelligence (Reinforcement Learning, Speedup Learning, Explanation-Based Learning).

Undergraduate Courses Taught:

Introduction to Artificial Intelligence Programming, Introduction to Theory of Computation, Social and Ethical Issues in Computing, Analysis of Algorithms, Introduction to Computer Science II (with C++ and Java).

Professional service:

- Action Editor, *Machine Learning*, 2006-current.
- Editorial Board Member, *Encyclopedia on Machine Learning*, Springer Verlag, 2006-current.
- Editorial Board Member, *Journal of Artificial Intelligence Research*, 1998-2001.
- Editorial Board Member, *Machine Learning*, 2004-2005.

- Local Arrangements Chair, International Conference on Machine Learning, 2007.
- Program Co-chair, International Conference on Inductive Logic Programming, 2007.
- Chair for the workshop on Relational Reinforcement Learning in International Machine Learning Conference in 2002.
- Area chair for the Reinforcement Learning area in International Machine Learning Conference in 2001, Machine Learning for National Conference on Artificial intelligence, 2009.
- Co-Chair of the workshop on “Knowledge Compilation and Speedup Learning” in International Machine Learning Conference in 1993. Chair of the workshop on “Knowledge Compilation and Speedup Learning” in International Machine Learning Conference in 1992.
- Tutorials:
 - “Reinforcement Learning: From Foundations to Advanced Techniques,” Prasad Tadepalli, Sridhar Mahadevan and Vivek Borkar at IJCAI 2007, Hyderabad, India.
 - “Decision-theoretic Planning and Reinforcement Learning in Relational Domains,” Prasad Tadepalli, Alan Fern and Kristian Kersting, at AAAI 2008, Chicago, IL, U.S.A.
- Senior program committee member: International Conference in Machine Learning, 2001, 2005; National Conference on Artificial Intelligence, 2005.
- Program committee member: International Conference on Machine Learning, 1995 – 2005, 2008, 2009; Inductive Logic Programming Conference, 1999 – 2003, 2007-2010; National Conference on Artificial Intelligence, 1990, 1994, 1997, 1998, 2000, 2005,2009; Neural Information Processing Systems, 2009; Workshop on Semantic Scientific Knowledge Integration at Spring Symposium on Artificial Intelligence, 2008; Brazilian Symposium on Artificial Intelligence, 2008; Knowledge Discovery and Data Mining, 2006; European Conference on Machine Learning, 2006; NIPS Workshop on Transfer Learning, 2005; KDD Workshop on Utility-Based Data Mining, 2005; ICML Workshop on Constrained Optimization and Structured Output Spaces, 2007; Indian Conference on Artificial Intelligence, 2003; FLAIRS Workshop on Learning in Planning and Scheduling, 2003; Algorithmic Learning Theory Conference, 1999; Computational Learning Theory conference, 1994; Workshop on Logic and Learning, 2001; ICML Workshop on learning text categorization, 1999; AAAI Fall Symposium on Planning and Learning in Games, 1993.
- Reviewer: National Science Foundation, Journal of Machine Learning Research, Machine Learning Journal, Artificial Intelligence Journal, Journal of Artificial Intelligence Research, IEEE Expert.
- Life Member of American Association of Artificial Intelligence (AAAI).

Departmental and University service:

- Member of University Graduate Council, 2003-2006, Interim member, 2007.
- Member of College on Engineering Graduate Curriculum Committee, 2009-2010.
- Head Graduate Advisor for Computer Science program 2001-current.
- Member of Masters in Engineering Program Committee, 2009.
- Chair of the comprehensive examination committee, 1993 Fall through 1997 spring, 1998 Fall through 2001 Spring.
- Member of graduate committee 1993-94, 1996-97, 1998-current.
- Member of undergraduate committee, 1999-00.
- Member of the equipment committee, 1991-92.
- Member of the curriculum committee, 1990-93.

Awards:

1. Work with Ronny Bjarnason and Alan Fern won student best paper award at the International Conference on Planning and Scheduling (ICAPS 2009).
2. Nominated for the College of Engineering Alumni Professor Award by the School of EECS, Oregon State University in 2008.

Grants:

1. "Partial Planning Reinforcement Learning," ARO, \$425,000, with Alan Fern, 2009-2012.
2. "ERUDITE: Machine Reading," DARPA, \$375,000/year with Tom Dietterich (PI) and Xiaoli Fern, 2009-2014.
3. "Integrated Learning," DARPA, \$433,181/year with Tom Dietterich (PI), Weng-Keen Wong, and Ron Metoyer, 2006-2008.
4. "CALO Architecture," DARPA, \$100,000 with Tom Dietterich (PI) and Alan Fern, 2005-2008.
5. "Effective Bayesian Transfer Learning," DARPA, \$800,000/year with Tom Dietterich (PI), Alan Fern, 2005-2008.
6. "Knowledge-Intensive Learning Methods," DARPA, \$472,259, with Tom Dietterich (PI), Bruce D'Ambrosio, and Jon Herlocker, 2003-04.
7. "Knowledge-Intensive Learning Methods," DARPA, \$511,000, with Tom Dietterich (PI), Alan Fern, and Jon Herlocker, 2004-05.
8. "Knowledge Plane: Technology Assessment," DARPA/USAF, \$129,853, with Tom Dietterich (PI), Alan Fern, and Ron Metoyer, 2005.
9. "Relational Reinforcement Learning," NSF, \$411,602, 2003-2006.
10. "Average Reward Reinforcement Learning:Scaling Up," NSF, \$368,659, 2001-04.
11. "Interdisciplinary Center for Computational Methods in Electronics Manufacturing," with Bruce D'Ambrosio (CS), Dean Jensen (IME), Rasaratnam Logendran (IME), and Robert Paasch (ME), Oregon State University Rickert Foundation seed grant, \$25,000, 1999-2000.
12. "Hybrid Computational Methods for Skill Acquisition," with Tom Dietterich (PI), ONR, \$379,200, 1995-98.
13. "Average Reward Reinforcement Learning," NSF, \$277,299, 1995-99.
14. "Tradeoffs in Learning and Planning," NSF, \$69,955, 1991-94.

Journal papers:

1. Natarajan, S., Tadepalli, P., Dietterich, T., and Fern, A. "Learning First-Order Probabilistic Models with Combining Rules" *Annals of Mathematics and Artificial Intelligence*, Special issue on Probabilistic Relational Learning, 2009.
2. Tadepalli, P. "Learning to Solve Problems from Exercises," *Computational Intelligence*, 24(4), 257-291, 2008.
3. Mehta, N., Natarajan, S., Tadepalli, P. and Fern, A., "Transfer in Variable-Reward Hierarchical Reinforcement Learning," *Machine Learning Journal*, 73(3), 289-312, 2008.
4. Dietterich, T. G., Domingos, P., Getoor, L., Muggleton, S. and Tadepalli, P. "Structured Machine Learning: The next 10 Years," *Machine Learning Journal*, 73:3-23, 2008.

5. Blockeel, H., Shaavlik, J. and Tadepalli, P., "Guest editors' introduction: special issue on inductive logic programming (ILP-2007), 73:1-2.
6. Bjarnason, R., Tadepalli, P. and Fern A., "Searching Solitaire in Real Time," *International Computer Games Association Journal*, 30:3,131-142,2007.
7. Tadepalli, P. "Cognitive Architectures have Limited Explanatory Power," Commentary on Anderson and Lebiere's article on "The Newell Test for a Theory of Mind," in *Behavioral and Brain Sciences*, 26(5), 2001.
8. Amoth, T., Cull, P. and Tadepalli, P. "On Exact Learning of Unordered Tree Patterns," *Machine Learning*, 44(3), 211-243, 2001.
9. Reddy, C. and Tadepalli, P. "Learning Horn Definitions: Theory and an Application to Planning," *New Generation Computing*, 17, 77-98, 1999.
10. Tadepalli, P. and Ok, D. "Model-based Average Reward Reinforcement Learning," *Artificial Intelligence*, 100, 177-224, 1998.
11. Tadepalli, P. and Russell, S., "Learning from Examples and Queries with Structured Determinations," *Machine Learning*, 32, 245-295, 1998.
12. Tadepalli, P. and Natarajan, B. "A Formal Framework for Speedup Learning from Problems and Solutions," *Journal of AI Research*, 4, 445-475, 1996.
13. Mahadevan, S. and Tadepalli, P., "Quantifying Prior Determination Knowledge using PAC Learning Model," *Machine Learning*, 17, 69-105, 1994.
14. Mahadevan, S., Mitchell, T., Mostow, J., Stienberg, L. and Tadepalli, P., "An Apprentice-based Approach to Knowledge Acquisition," *AI Journal*, 64, 1-52, 1993.
15. Weiss, S., Galen, R., and Tadepalli, P., "Maximizing the Predictive Value of Production Rules," *AI Journal*, 45, 47-71, 1990.
16. Hall, P.R., Falkenhainer, B., Flann, N. S., Hampson, S., Reinke, R., Shrager, J., Sims, M., and Tadepalli, P., "A review of the fourth international workshop on machine learning," *Machine Learning*, 2 (2), 1987.

Pending Journal Papers

1. Natarajan, S., Judah, K., Fern, A., and Tadepalli, P. "A Decision-Theoretic Model of Assistance," submitted to *Journal of Artificial Intelligence Research*.
2. Parker, C., Fern, A. and Tadepalli, P., "Gradient Boosting for Efficient and Accurate Retrieval of Structured Sequences," submitted to *Journal of Artificial Intelligence Research*.

Book Chapters and Edited Collections:

1. Shavlik, J., Blockkeel, H., Ramon, J. and Tadepalli, P., (Ed), *Proceedings of The 17th International Conference on Inductive Logic Programming*, Springer Verlag, 2007.
2. Mehta, N., Natarajan, S., Tadepalli, P. and Fern, A., "Transfer in Hierarchical Variable Reward Reinforcement Learning," (accepted), in *Transfer Learning*, Silver, D., Caruana, R. and Bennett, C., (Ed.), Springer Verlag, 2007.
3. Tadepalli, P., "Learning in Intractable Domains," in *Machine Learning: A Guide to Current Research*, Mitchell, T. M., Michalski, R. S., and Carbonell, J. G., (Ed.), Morgan Kaufmann, 1986.

Rigorously refereed papers in highly selective conferences:

1. Bjarnason, R., Fern, A. and Tadepalli, P., "Lower Bounding Klondike Solitaire with Monte-Carlo Planning," in International Conference on Automated Planning and Scheduling (ICAPS), 2009.
2. Proper, S., Tadepalli, P., "Solving Multiagent Assignment Markov Decision Processes," in Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems, 2009.
3. Proper, S., Tadepalli, P., "Transfer Learning via Relational Templates," in Proceedings of the 19th International Conference on Inductive Logic Programming, 2009.
4. Bjarnason, R., Tadepalli, P., Fern, A. and Niedner, C., "Simulation-based Optimization of Resource Placement and Emergency Response," in Conference on Innovative Applications of Artificial Intelligence (IAAI), 2009.
5. Mehta, N., Ray, S., Tadepalli, P., and Dietterich, T. "Automatic Discovery and Transfer of MAXQ Hierarchies," in International Conference on Machine Learning, 2008.
6. Natarajan, S., Bui, H., Tadepalli, P., Kersting, K., Wong, W-K. "Logical Hierarchical Hidden Markov Models for User Activity Recognition," in International Conference on Inductive Logic Programming, 2008.
7. Natarajan, S., Tadepalli, P., and Fern A. "A Relational Hierarchical Model for Decision Theoretic Assistance," in International Conference on Inductive Logic Programming, 2007.
8. Parker, C., Fern, A., and Tadepalli, P., "Learning for Efficient Retrieval of Structured Data with Noisy Queries," in International Conference on Machine Learning, 2007.
9. Wilson, A., Fern, A., Ray, S. and Tadepalli, P. "Multi-Task Reinforcement Learning: A Hierarchical Bayesian Approach," in International Conference on Machine Learning, 2007.
10. Fern, A., Natarajan, S., Judah, K., and Tadepalli, P., "A Decision-theoretic Model of Assistance," in International Joint Conference on Artificial Intelligence, 2007.
11. Proper, S. and Tadepalli, P., "Scaling Model-Based Average-Reward Reinforcement Learning for Product Delivery," in European Conference on Machine Learning, 2006.
12. Parker, C., Fern, A. and Tadepalli, P. "Gradient Boosting for Sequence Alignment," in National Conference on Artificial Intelligence, 2006.
13. Natarajan, S. and Tadepalli, P. "Dynamic Preferences in Multi-criteria Reinforcement Learning," in International Conference on Machine Learning, 2005.
14. Natarajan, S., Tadepalli, P., Altendorf, E., Dietterich, T., Fern, A. and Restificar, A., "Learning First-Order Probabilistic Models with Combining Rules," in International Conference on Machine Learning, 2005.
15. Chisholm, M. and Tadepalli, P. "Learning Decision Rules from Randomized Iterative Local Search" in Proceedings of the International Conference on Machine Learning, 2002.
16. Seri, S. and Tadepalli, P. "Model-based Hierarchical Average-Reward Reinforcement Learning," in Proceedings of the International Conference on Machine Learning, 2002.
17. Amoth, T., Cull, P. and Tadepalli, P. "Exact Learning of Unordered Tree Patterns from Queries," in Proceedings of the Conference on Computational Learning Theory, 1999.
18. Amoth, T., Cull, P. and Tadepalli, P. "Exact Learning of Tree Patterns from Queries and Counterexamples," in Proceedings of the Conference on Computational Learning Theory, 1998.

19. Reddy, C. and Tadepalli, P., "Learning First-Order Acyclic Horn Programs from Entailment," in Proceedings of International Conference on Machine Learning, 1998, also appeared in International Conference on Inductive Logic Programming, 1998.
20. Tadepalli, P. and Dietterich, T. G. "Hierarchical Explanation-Based Reinforcement Learning," in Proceedings of International Machine Learning Conference, 1997.
21. Liere, R. and Tadepalli, P. "Active Learning with Committees for Text Categorization," in Proceedings of National Conference on Artificial Intelligence, 1997.
22. Reddy, C., Tadepalli, P. "Learning Goal-Decomposition Rules using Exercises," in Proceedings of International Conference on Machine Learning, 1997.
23. Tadepalli, P. and Ok, D. "Scaling up Average Reward Reinforcement Learning by Approximating the Domain Models and the Value Function," in Proceedings of International Machine Learning Conference, 1996.
24. Ok, D. and Tadepalli, P., "Auto-exploratory Average Reward Reinforcement Learning," in Proceedings of National Conference on Artificial Intelligence, 1996.
25. Reddy, C., Tadepalli, P. and Roncagliolo, S., "Theory-guided Empirical Speedup Learning of Goal Decomposition Rules," in Proceedings of International Machine Learning Conference, 1996.
26. Tadepalli, P., "Learning from Queries and Examples with Tree-Structured Bias," in Proceedings of International Machine Learning Conference, Amherst, MA, 1993.
27. Tadepalli, P., "A Theory of Unsupervised Speedup Learning," in Proceedings of National Conference on Artificial Intelligence, San Jose, CA, 1992.
28. Tadepalli, P., "A Formalization of Explanation-Based Macro-operator Learning," in International Joint Conference on Artificial Intelligence, Sydney, Australia, 1991.
29. Tadepalli, P., "Lazy Explanation-Based Learning: A Solution to the Intractable Theory Problem," Proceedings of International Joint Conference on Artificial Intelligence, Detroit, MI, 1989.
30. Natarajan, B. and Tadepalli, P., "Two New Frameworks for Learning," in Proceedings of the International Machine Learning Conference, Ann Arbor, MI, 1988.
31. Mahadevan, S. and Tadepalli, P., "On the Tractability of Learning from Incomplete Theories," in Proceedings of the International Machine Learning Conference, Ann Arbor, MI, 1988.
32. Weiss, S., Galen, R., and Tadepalli, P., "Optimizing the Predictive Value of Diagnostic Decision Rules," in Proceedings of the National Conference of AAAI-87, Seattle, 1987.

Refereed workshop/symposium papers:

1. Natarajan, S., Tadepalli, P., Kunapuli, G. and Shavlik, J. "Knowledge Intensive Learning: Directed vs. Undirected SRL Models," in Workshop on Statistical Relational Learning, 2009.
2. Mehta, N., Ray, S., Tadepalli, P., and Dietterich, T. "Automatic Discovery and Transfer of MAXQ Hierarchies," in NIPS Workshop on Hierarchical Organization of Behavior: Computational, Psychological and Neural Perspectives, 2007.
3. Parker, C., Tadepalli, P., Wong, W-K., Dietterich, T., and Fern, A., "Learning from Demonstrations via Structured Prediction," in AAAI Workshop on Acquiring Planning Knowledge via Demonstrations, 2007.
4. Pierce, C-M., Wong, W-K., Tadepalli, P., and Dereszynski, E., "Bi-level Optimization for Learning Cost Functions from Demonstration," in AAAI Workshop on Acquiring Planning Knowledge via Demonstrations, 2007.

5. Natarajan, S., Judah, K., Tadepalli, P. and Fern, A., "A Decision-Theoretic Model of Assistance - Evaluation, Open Problems and Extensions," AAAI Spring Symposium on Interaction Challenges for Intelligent Assistants, 2007.
6. Natarajan, S., Tadepalli, P. and Fern, A., "Exploiting prior Knowledge in Intelligent Assistants - Combining relational models with hierarchies," Extended Abstract in the Proceedings of the Dagstuhl Seminar on Probabilistic, Logical and Relational Learning - A Further Synthesis, 2007.
7. Fern, A., Natarajan, S., Judah, K., and Tadepalli, P., "A Decision-Theoretic Model of Assistance," AAAI Workshop on Modeling Others from Observation, 2006.
8. Natarajan, S., Wong, W-K., and Tadepalli, P., "Structure Refinement in First-Order Conditional Influence Language," ICML Workshop on Statistical Relational Learning, 2006.
9. Mehta, N., Natarajan, S., Tadepalli, P., and Fern, A., "Transfer in Hierarchical Variable-Reward Reinforcement Learning," NIPS Workshop on Inductive Transfer, 2005.
10. Mehta, N. and Tadepalli, P., "Multi-agent Shared Hierarchy Reinforcement Learning," Workshop on Rich Representations for Reinforcement Learning, Bonn, Germany, 2005.
11. Proper, S. and Tadepalli, P. "Scaling Average-reward Reinforcement Learning for Product Delivery," in Proceedings of the Fall Symposium on Artificial Intelligence, Washington, D.C. 2004.
12. Proper, S., Tadepalli, P., Tang, H., and Logendran, R. "A Reinforcement Learning Approach to Multiple Vehicle Product Delivery," in Proceedings of the International Conference on Industrial Engineering, Portland, OR, 2003.
13. Tadepalli, P., Givan, R. and Driessens, K. "Relational Reinforcement Learning: An Overview," in Proceedings of the Workshop on Relational Reinforcement Learning, at International Conference on Machine Learning, Banff, Canada, 2004.
14. Roncagliolo, S. and Tadepalli, P., "Function Approximation in Hierarchical Relational Reinforcement Learning," in Proceedings of the Workshop on Relational Reinforcement Learning, at International Conference on Machine Learning, Banff, Canada, 2004.
15. Liere, R. and Tadepalli, P. "The Use of Active Learning in Text Categorization," In AAAI Spring Symposium on Machine Learning in Information Access, 1996.
16. Tadepalli, P. and Ok, D., "Discounting Considered Harmful: A Comparison of Reinforcement Learning Techniques in Automatic Guided Vehical Scheduling," in Proceedings of the Robot Learning workshop, New Brunswick, NJ, 1994.
17. Tadepalli, P., Isukapalli, R. and Roncagliolo, S., "Speedup Learning as Exploiting Problem Space Structure" in Proceedings of the workshop on Knowledge Compilation and Speedup Learning, Amherst, MA, 1993.
18. Roncagliolo, S. and Tadepalli, P., "Empirical Speedup Learning of Decomposition Rules for Planning," in Proceedings of the workshop on Knowledge Compilation and Speedup Learning, Amherst, MA, 1993.
19. Tadepalli, P. and Isukapalli, R., "Learning Plan Knowledge Using Simulators," in Proceedings of the workshop on Knowledge Compilation and Speedup Learning, Aberdeen, Scotland, 1992.
20. Tadepalli, P. and Joshi, V., "Realtime Scheduling using Minimin Search," in Proceedings of AAAI Spring Symposium Series on Practical Approaches to Planning and Scheduling, Stanford, CA, 1992.
21. Tadepalli, P., "Learning with Inscrutable Theories," in International Machine Learning Workshop, Chicago, IL, 1991.

22. Tadepalli, P., "On Quantifying Approximation," Proceedings of the Workshop on Automatic Generation of Abstractions and Approximations at AAAI National Conference, Boston, MA, 1990.
23. Tadepalli, P., "Lazy Explanation-Based Learning," Proceedings of the Workshop on Explanation at AAAI National Conference, Boston, MA, 1990.
24. Tadepalli, P., "Planning in Games Using Approximately Learned Macros," Proceedings of Machine Learning Workshop, Ithaca, NY, 1989.
25. Mahadevan, S., Natarajan, B., and Tadepalli, P., "A Framework for Learning as Improving Problem-Solving Performance," in Proceedings of the Spring Symposium Series on Explanation-Based Learning, Stanford University, Palo Alto, CA, 1988.

Technical reports and other publications:

1. Tadepalli, P., Etzioni, O., Fisher, D., Flann, N., Minton, S., Prieditis, A., Subramanian, D., van Hermelen, F., (Ed.), "Proceedings of the Workshop on Knowledge Compilation and Speedup Learning," Oregon State University Technical Report, 92-30-7.
2. Fisher, D., Subramanian, D., and Tadepalli, P., "An Overview of Current Research on Knowledge Compilation and Speedup Learning" in Proceedings of the workshop on Knowledge Compilation and Speedup Learning, Aberdeen, Scotland, 1992.
3. Yenong, Q. and Tadepalli, P., "Local search Methods for Job-Shop Scheduling," Oregon State University Technical Report, 94-30-2.
4. Tadepalli, P. and Ok, D., "H-Learning: A Reinforcement Learning Method to Optimize Undiscounted Average Reward," Oregon State University Technical Report, 94-30-1.
5. Dietterich, T., Ok, D., Tadepalli, P. and Zhang, W., "Reinforcement Learning in Scheduling," in Space Operations and Applications Conference, Houston, TX, 1993.
6. Tadepalli, P., "Knowledge-Based Planning in Games," Carnegie Mellon University Technical Report, CMU-CS-89-135, 1989.
7. Tadepalli, P., "Learning Approximate Plans in Games," Ph.D. Thesis, Rutgers University Technical Report, ML-TR-8, 1986.

Students graduated:

1. Sriraam Natarajan, Ph.D., 2007.
Thesis: Effective Decision-Theoretic Assistance Through Relational Hierarchical Models.
2. Charles Parker, Ph.D., 2007.
Thesis: Structured Gradient Boosting.
3. Kiran Polavarapu, M.S., 2006.
Project: Event and Sentiment Extraction in the Financial Domain
4. Thierry Donneaugolencer, M.S., 2005
Project: Planning in Multi-agent Partially Observable Domains using Sparse Sampling.
5. Sriraam Natarajan, M.S., 2004.
Thesis: Multi-criteria Average-reward Reinforcement Learning.
6. Nimish Dharawat (M.S.), 2004
Project: Learning Tree Patterns for Information Extraction.
7. Kim Mach, M.S., 2004.
Project: An Experimental Evaluation of Auto-exploratory Average-reward Reinforcement Learning.

8. Hong Tang, M.S., 2002.
Project: Average-reward Reinforcement Learning for Product Delivery by Multiple Vehicles.
9. Sandeep Seri, M.S., 2002.
Thesis: Hierarchical Average-reward Reinforcement Learning.
10. Tom Amoth, Ph.D., 2001. Co-advisor with Prof. Paul Cull.
Thesis: Exact Learning of Tree Patterns. Nominated by the department for the ACM distinguished thesis award.
11. Michael Chisholm, M.S., 1999.
Thesis: Learning Classification Rules by Randomized Iterative Local Search
12. Ray Liere, Ph.D., 1999.
Thesis: Active Learning with Committees: An Approach To Efficient Learning in Text Categorization Using Linear Threshold Algorithms
13. Chandra Reddy, Ph.D., 1998.
Thesis: Learning Hierarchical Decomposition Rules for Planning: An Inductive Logic Programming Approach.
14. Dokyeong Ok, Ph.D., 1996.
Thesis: A Study of Model-based Average Reward Reinforcement Learning.
15. Peter Dudey, M.S., 1995.
Thesis: Constructive Induction for Improved Learning of Boolean Functions.
16. Qi Yenong, M.S., 1993.
Project: Local Search Methods for Job-Shop Scheduling.
17. Silvana Roncagliolo, M.S., 1993.
Project: Empirical Speedup Learning of Decomposition Rules for Planning.
18. Ramana Isukapalli, M.S., 1992.
Thesis: Learning Macro-operators for Planning using Simulators.

References

Available upon request.