

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 (1984-90) 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), Ethics and Diversity in EECS.

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.
- Associate Editor, *Journal of Artificial Intelligence Research*, 2010-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.
- Co-chair of senior member track, National Conference on Artificial Intelligence, 2017-2018.
- Senior program committee member/area chair: International Conference in Machine Learning, 2001, 2005; National Conference on Artificial Intelligence, 2005,2011,2012,2013,2015,2016,2017,2017,2018,2019; International Joint Conference on Artificial Intelligence, 2011,2016,2017, 2018;
- Program committee member: International Conference on Parallel Processing, 2018; Annual Meeting of Association of Computational Linguistics, 2018; International Conference on Machine Learning, 1995 – 2005, 2008, 2009, 2011, 2012,2013,2018; Inductive Logic Programming Conference, 1999 – 2003, 2007-2010, 2012; National Conference on Artificial Intelligence, 1990, 1994, 1997, 1998, 2000, 2005,2009; International Conference on Planning and Scheduling, 2011, 2012, 2013, 2016; Workshop on Internattional Planning Competition, 2012. Workshop on Statistical Relational Learning, 2012. Workshop on Statistical Relational Artificial Intelligence, 2012, 2013. Neural Information Processing Systems, 2009, 2010, 2011, 2012; 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, 2011, 2012; 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).
- Member of Association of Computing Machinery (ACM).

Departmental and University service:

- Member of University Graduate Council, 2003-2006, Interim member, 2007.
- Member of College of Engineering Change team, 2016-2017.
- Member of College on Engineering Graduate Curriculum Committee, 2009-2010.

- Head Graduate Advisor for Computer Science program 1998-current.
- Member of Promotion and Tenure Dossier Committee, 2010-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. Outstanding paper award, AAAI 2013, with Janardhan Rao Doppa and Alan Fern.

Grants:

1. "EXACT: Explanation-Informed Acceptance Testing of Deep Adaptive Programs," with Alan Fern (PI), Tom Dietterich, Margaret Burnett, Martin Erwig, Liang Huang, and Fuxin Li, \$6,545,123, 2017-2021.
2. "RI: Small: Speedup Learning for Online Planning Under Uncertainty," NSF, with Alan Fern (PI), \$450,000, 2016-2019.
3. "Active Transfer Learning of Latent Competencies," ONR, STTR Phase 1 with Robby Robson (PI) of Eduworks, \$25K 2015-2016.
4. "II-SN: Software Tools for Monte-Carlo Optimization," NSF, \$442,366, with Alan Fern (PI), Alex Groce, Sinisa Todorovic, and Tom Dietterich 2014-2017.
5. "RI:Small: Integrating Learning and Search for Structured Prediction," NSF, \$449,997, with Alan Fern, 2012-2015
6. "Deep Reading and Learning," DARPA, \$2,221,485, with Thomas G. Dietterich and Xiaoli Fern, 2012-2016
7. "First-Order Decision Theoretic Planning," NSF, \$267,500, CI fellows award for Saket Joshi, 2010-2012
8. "Lifelong Active Transfer Learning for Sequential Decision Making," ONR, \$1,150,000, with Alan Fern, 2011-2015
9. "RI:Medium: Optimizing Policies for Service Organizations in Complex Structured Domains," NSF, \$581,680, with Alan Fern, 2010-2014
10. "A Compute Cluster and Software Tools for Monte-Carlo Methods in Artificial Intelligence, " NSF, \$600,000, with Tom Dietterich (PI), Alan Fern, Kagan Tumer, and Weng-Keen Wong, 2010-2014
11. "Partial Planning Reinforcement Learning," ARO, \$425,000, with Alan Fern, 2009-2012.
12. "ERUDITE: Machine Reading," DARPA, \$375,000/year with Tom Dietterich (PI) and Xiaoli Fern, 2009-2012.
13. "Integrated Learning," DARPA, \$433,181/year with Tom Dietterich (PI), Weng-Keen Wong, and Ron Metoyer, 2006-2008.

14. “CALO Architecture,” DARPA, \$100,000 with Tom Dietterich (PI) and Alan Fern, 2005-2008.
15. “Effective Bayesian Transfer Learning,” DARPA, \$800,000/year with Tom Dietterich (PI), Alan Fern, 2005-2008.
16. “Knowledge-Intensive Learning Methods,” DARPA, \$472,259, with Tom Dietterich (PI), Bruce D’Ambrosio, and Jon Herlocker, 2003-04.
17. “Knowledge-Intensive Learning Methods,” DARPA, \$511,000, with Tom Dietterich (PI), Alan Fern, and Jon Herlocker, 2004-05.
18. “Knowledge Plane: Technology Assessment,” DARPA/USAF, \$129,853, with Tom Dietterich (PI), Alan Fern, and Ron Metoyer, 2005.
19. “Relational Reinforcement Learning,” NSF, \$411,602, 2003-2006.
20. “Average Reward Reinforcement Learning:Scaling Up,” NSF, \$368,659, 2001-04.
21. “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.
22. “Hybrid Computational Methods for Skill Acquisition,” with Tom Dietterich (PI), ONR, \$379,200, 1995-98.
23. “Average Reward Reinforcement Learning,” NSF, \$277,299, 1995-99.
24. “Tradeoffs in Learning and Planning,” NSF, \$69,955, 1991-94.

Journal papers:

1. Fern, A., Natarajan, S., Judah, K., and Tadepalli, P., “A Decision-Theoretic Model of Assistance,” *Journal Of Artificial Intelligence Research* 50:71–104, 2014.
2. Judah, K., Fern, A., Dietterich, T. G., and Tadepalli, P., “Active Imitation Learning: Formal and Practical Reductions to I.I.D. Learning,” *Journal of Machine Learning Research*, 15: 3925–3963, 2014.
3. Doppa, J., Fern, A., and Tadepalli, P., “HC-Search: A Learning Framework for Search-based Structured Prediction,” *Journal of Artificial Intelligence Research*, (JAIR), 50:369–407, 2014.
4. Wilson, A., Fern, A., and Tadepalli, P., “Using Trajectory Data to Improve Bayesian Optimization for Reinforcement Learning,” *Journal of Machine Learning Research (JMLR)*, 15, 253–282, 2014.
5. Doppa, J., Fern, A., and Tadepalli, P., “Structured Prediction via Output Space Search,” *Journal of Machine Learning Research*, (JMLR), 15(1), 1317–1350, 2014.
6. Natarajan, S., Tadepalli, P., and Fern, A., “A Relational Hierarchical Model of Decision-Theoretic Assistance,” *Knowledge and Information Systems(KAIS)*, 32(2), 329-349, 2012.
7. Zhang, X. Yoon, S. W., DiBona, P., etal, “An Ensemble Architecture for Learning Complex Problem Solving Techniques from Demonstration,” *ACM Transactions on Intelligent Systems and Technology* , 3(4), 75, 2012.
8. Wilson, A., Fern, A., Tadepalli, P., “Transfer Learning in Sequential Decision Problems: A Hierarchical Bayesian Approach,” *Journal of Machine Learning Research - Proceedings Track*, 27: 217-227, 2012.
9. Fern, A., Khardon, R., Tadepalli, P., “The First Learning Track of the International Planning Competition,” *Machine Learning*, 84:81-107, 2011.
10. Mehta, N., Ray, S., Tadepalli, P., and Dietterich, T., “Automatic Discovery and Transfer of Task Hierarchies in Reinforcement Learning,” *AI Magazine*, 32(1), 2011.

11. Parker, C., Altun, Y., Tadepalli, P. “Guest Editorial: Special Issue on Structured Prediction,” *Machine Learning*, 77(2-3), September 2009.
12. 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, 54(1-3), 2008.
13. Tadepalli, P. “Learning to Solve Problems from Exercises,” *Computational Intelligence*, 24(4), 257–291, 2008.
14. Mehta, N., Natarajan, S., Tadepalli, P. and Fern, A., “Transfer in Variable-Reward Hierarchical Reinforcement Learning,” *Machine Learning Journal*, 73(3), 289–312, 2008.
15. 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.
16. Blockeel, H., Shaavlik, J. and Tadepalli, P., “Guest editors’ introduction: special issue on inductive logic programming (ILP-2007), 73:1-2.
17. Bjarnason, R., Tadepalli, P. and Fern A., “Searching Solitaire in Real Time,” *International Computer Games Association Journal*, 30:3,131–142,2007.
18. 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.
19. Amoth, T., Cull, P. and Tadepalli, P. “On Exact Learning of Unordered Tree Patterns,” *Machine Learning*, 44(3), 211–243, 2001.
20. Reddy, C. and Tadepalli, P. “Learning Horn Definitions: Theory and an Application to Planning,” *New Generation Computing*, 17, 77–98, 1999.
21. Tadepalli, P. and Ok, D. “Model-based Average Reward Reinforcement Learning,” *Artificial Intelligence*, 100, 177–224, 1998.
22. Tadepalli, P. and Russell, S., “Learning from Examples and Queries with Structured Determinations,” *Machine Learning*, 32, 245–295, 1998.
23. Tadepalli, P. and Natarajan, B. “A Formal Framework for Speedup Learning from Problems and Solutions,” *Journal of AI Research*, 4, 445-475, 1996.
24. Mahadevan, S. and Tadepalli, P., “Quantifying Prior Determination Knowledge using PAC Learning Model,” *Machine Learning*, 17, 69–105, 1994.
25. Mahadevan, S., Mitchell, T., Mostow, J., Stienberg, L. and Tadepalli, P., “An Apprentice-based Approach to Knowledge Acquisition,” *AI Journal*, 64, 1-52, 1993.
26. Weiss, S., Galen, R., and Tadepalli, P., “Maximizing the Predictive Value of Production Rules,” *AI Journal*, 45, 47-71, 1990.
27. 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.

Book Chapters and Edited Collections:

1. Tadepalli, P., “Average-Reward Reinforcement Learning,” in *Encyclopedia of Machine Learning 2010*, 64-68.
2. Ray, S. and Tadepalli, P., “Model-Based Reinforcement Learning,” in *Encyclopedia of Machine Learning 2010*, 690-693.

3. Shavlik, J., Blockeel, H., Ramon, J. and Tadepalli, P., (Ed), Proceedings of The 17th International Conference on Inductive Logic Programming, Springer Verlag, 2007.
4. Mehta, N., Natarajan, S., Tadepalli, P. and Fern, A., “Transfer in Hierarchical Variable Reward Reinforcement Learning,” in *Transfer Learning*, Silver, D., Caruana, R. and Bennett, C., (Ed.), Springer Verlag, 2007.
5. 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 conferences:

1. Dayapule, D. H. Raghavan, A., Tadepalli, P., Fern, A., “Emergency Response Optimization using Online Hybrid Planning,” IJCAI 2018.
2. Shahbazi, H., Fern, X. Z., Ghaeini, R., Tadepalli, P., “Joint Neural Entity Linking with Output Space Search,” COLING, 2018.
3. Ghaeini, R., Fern, X. Z., Shahbazi, H., Tadepalli, P. “Dependent Gated Reading for Cloze-Style Question Answering,” COLING, 2018.
4. Fern, A. Goetschalckx, R., Hamidi-Haines, M., Tadepalli, P. “Adaptive Submodularity with Varying Query Sets: An Application to Active Multi-label Learning,” Proceedings of Machine Learning Research (PMLR), Vol 76, pp 577-592, 2017.
5. Ma, C. Doppa, J. R., Tadepalli, P. Shahbazi, H., Fern, X., “Multi-Task Structured Prediction for Entity Analysis: Search-based Learning Algorithms,” Proceedings of Machine Learning Research (PMLR), Vol 77, pp 514-529, 2017.
6. Raghavan, A., Sanner, S., Khardon, R., Tadepalli, P., Fern, A., “Hindsight Optimization for Hybrid State and Action MDPs,” AACL, 2017, 3790-3796.
7. Obeidat, R., Fern, X. Z., Tadepalli, P., “Label Embedding Approach for Transfer Learning,” ICBO/BioCreative, 2016.
8. Ghaeini, R., Fern, X., Huang, L., Tadepalli, P., “Event Nugget Detection with Forward-Backward Recurrent Neural Networks,” ACL (2) 2016.
9. Issakkimuthu, M., Fern, A., Khardon, R., Tadepalli, P., Xue, S., “Hindsight Optimization for Probabilistic Planning with Factored Actions,” ICAPS, 2015, 120-128.
10. Hamidi, M., Tadepalli, P., Goetschalckx, R., and Fern, A., “Active Imitation Learning of Hierarchical Policies,” in Proceedings of International Joint Conference on Artificial Intelligence, IJCAI, 2015.
11. Goetschalckx, R., Fern, A., and Tadepalli, P., “Multi-task Coactive Learning,” in Proceedings of International Joint Conference on Artificial Intelligence, IJCAI, 2015.
12. Xie, J., Ma, C., Doppa, J. R., Mannem, P., Fern, X. Z., Dietterich, T. G., and Tadepalli, P., “Learning Greedy Policies for the Easy-First Framework,” National Conference on Artificial Intelligence, NCAI, 2339-2345, 2015.
13. Raghavan, A., Khardon, R., Tadepalli, P., and Fern, A., “Memory-Efficient Symbolic Online Planning for Factored MDPs,” UAI, 732-741, 2015.
14. Cui, H., Khardon, R., Fern, A., and Tadepalli, P., “Factored MCTS for Large Scale Stochastic Planning,” National Conference on Artificial Intelligence, NCAI, 3261-3267, 2015.
15. Orr, J. W., Tadepalli, P., Doppa, J. R., Fern, X., and Dietterich, T., “Learning Scripts as Hidden Markov Models,” Proceedings of the National Conference on Artificial Intelligence, NCAI, 2014.

16. Judah, K., Fern, A., Tadepalli, P., and Goetschalckx, R., “Imitation Learning with Demonstrations and Shaping Rewards,” Proceedings of the National Conference on Artificial Intelligence, NCAI, 2014.
17. Goetschalckx, R., Fern, A., and Tadepalli, P., “Coactive Learning for Locally Optimal Problem Solving,” Proceedings of the National Conference on Artificial Intelligence, NCAI, 2014.
18. Doppa, J. R., Yu, J., Ma, C., Fern, A., and Tadepalli, P., “HC-Search for Multi-label Prediction: An Empirical Study,” Proceedings of National Conference on Artificial Intelligence, 2014.
19. Ma, C., Doppa, J. R., Orr, J. W., Mannem, P., Fern, X. Z., Dietterich, T. G., Tadepalli, P., “Prune-and-Score: Learning for Greedy Coreference Resolution,” Proceedings of the Conference on Empirical Methods in Natural Language Processing, 2115-2126, 2014.
20. Raghavan, A., Fern, A., Tadepalli, P., and Khardon, R., “Symbolic Opportunistic Policy Iteration for Factored-Action MDPs,” Proceedings of the International Conference on Neural Information Processing Systems (NIPS), 2499–2507, 2013.
21. Joshi, S., Khardon, R., Tadepalli, P., Raghavan, A., and Fern, A. “Solving Relational MDPs with Exogenous Events and Additive Rewards,” European Conference on Machine Learning, 178–193, 2013.
22. Natarajan, S., Odom, P., Joshi, S., Khot, T., Kersting, K., and Tadepalli, P., “Accelerating Imitation Learning in Relational Domains via Transfer by Initialization,” Proceedings of International Conference on Inductive Logic Programming, 2013.
23. Doppa, J. R., Fern, A., and Tadepalli, P., “HC-Search: Learning Heuristics and Cost Functions for Structured Prediction,” Proceedings of National Conference on Artificial Intelligence, 2013.
Outstanding Paper Award.
24. Raghavan, A., Joshi, S., Fern, A., Tadepalli, P. and Khardon, R., “Planning in Factored Action Spaces using Symbolic Dynamic Programming,” Proceedings of National Conference on Artificial Intelligence, 2012.
25. Doppa, J., Fern, A. and Tadepalli, P., “Output Space Search for Structured Prediction,” Proceedings of International Conference on Machine Learning, 2012.
26. Wilson, A., Fern, A., and Tadepalli, P., “A Bayesian Approach for Policy Learning from Trajectory Preference Queries,” Neural Information Processing Systems, 2012, 1142-1150
27. Sorower, S., Dietterich, T. G., Doppa, J. R., Orr, W., Tadepalli, P. and Fern, X.: Inverting Grice’s Maxims to Learn Rules from Natural Language Extractions, In Neural Information Processing Systems, 2011, 1053-1061.
28. Mehta, N., Tadepalli, P., and Fern, A. Autonomous Learning of Action Models for Planning. In Neural Information Processing Systems, 2011, 2465-2473.
29. Natarajan, S., Joshi, S., Tadepalli, P., Kersting, K., and Shavlik, J. W., “Imitation Learning in Relational Domains: A Functional Gradient Boosting Approach,” In International Joint Conference on Artificial Intelligence, 2011, 1414-1420.
30. Doppa, J. R., Sorower, M. S., NasrEsfahani, M., Orr, W., Dietterich, T. G., Fern, X., Tadepalli, P., and Irvine, J., “Learning Rules from Incomplete Examples via Implicit Mention Models,” In Journal of Machine Learning Research - Proceedings Track 20: 197-212 (2011).
31. Fern, A. and Tadepalli, P., “A Computational Decision Theory for Interactive Assistants,” in Neural Information Processing Systems, 2010, 577-585.

32. Natarajan, S., Khot, T., Lowd, D., Tadepalli, P., Kersting, K., and Shavlik, J. W., "Exploiting Causal Independence in Markov Logic Networks: Combining Undirected and Directed Models," ECML/PKDD, 2010, 434-450
33. Natarajan, S., Kunapuli, G., Judah, J., Tadepalli, P., Kersting, K., and Shavlik, J. W., "Multi-Agent Inverse Reinforcement Learning," in International Conference on Machine Learning and Applications, 2010, 395-400.
34. Wilson, A., Fern, A., and Tadepalli, P., "Bayesian Policy Search for Multi-Agent Role Discovery," in National Conference on Artificial Intelligence, 2010
35. Wilson, A., Fern, A., and Tadepalli, P., "Bayesian Role Discovery for Multi-agent Reinforcement Learning," in Autonomous Agents and Multi-agent Systems, 2010, 1587-1588
36. Wilson, A., Fern, A., and Tadepalli, P., "Incorporating Domain Models into Bayesian Optimization for Reinforcement Learning," in European Conference on Machine Learning, 2010, 467-482
37. Doppa, J. R., Yu, J., Tadepalli, P., and Getoor, L., "Learning Algorithms for Link Prediction Based on Chance Constraints," in European Conference on Machine Learning, 2010, 344-360
38. 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, **Best Student Paper Award**.
39. Proper, S., Tadepalli, P., "Solving Multiagent Assignment Markov Decision Processes," in Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems, 2009.
40. Proper, S., Tadepalli, P., "Transfer Learning via Relational Templates," in Proceedings of the 19th International Conference on Inductive Logic Programming, 2009.
41. Natarajan, S., Tadepalli, P., Kunapuli, G., and Shavlik, J., "Learning Parameters for Relational Probabilistic Models with Noisy-Or Combining Rule," in Proceedings of the International Conference on Machine Learning and Applications, 2009.
42. Bjarnason, R., Tadepalli, P., Fern, A. and Niedner, C., "Simulation-based Optimization of Resource Placement and Emergency Response," in Proceedings of the Conference on Innovative Applications of Artificial Intelligence (IAAI), 2009.
43. Proper, S. and Tadepalli, P., "Multiagent Transfer Learning via Assignment-based Decomposition," in Proceedings of the International Conference on Machine Learning and Applications, 345-350, 2009.
44. Zhang, X., Yoon, S., Dibona, P., Appling, D., Ding, L., Doppa, J., Green, D., et al., "An Ensemble Learning and Problem Solving Architecture for Airspace Management," in Proceedings of the Conference on Innovative Applications of Artificial Intelligence (IAAI), 2009.
45. Mehta, N., Ray, S., Tadepalli, P., and Dietterich, T. "Automatic Discovery and Transfer of MAXQ Hierarchies," in International Conference on Machine Learning, 2008.
46. 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.
47. Natarajan, S., Tadepalli, P., and Fern A. "A Relational Hierarchical Model for Decision Theoretic Assistance," in International Conference on Inductive Logic Programming, 2007.
48. Parker, C., Fern, A., and Tadepalli, P., "Learning for Efficient Retrieval of Structured Data with Noisy Queries," in International Conference on Machine Learning, 2007.

49. Wilson, A., Fern, A., Ray, S. and Tadepalli, P. “Multi-Task Reinforcement Learning: A Hierarchical Bayesian Approach,” in International Conference on Machine Learning, 2007.
50. Proper, S. and Tadepalli, P., “Scaling Model-Based Average-Reward Reinforcement Learning for Product Delivery,” in European Conference on Machine Learning, 2006.
51. Parker, C., Fern, A. and Tadepalli, P. “Gradient Boosting for Sequence Alignment,” in National Conference on Artificial Intelligence, 2006.
52. Natarajan, S. and Tadepalli, P. “Dynamic Preferences in Multi-criteria Reinforcement Learning,” in International Conference on Machine Learning, 2005.
53. 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.
54. Chisholm, M. and Tadepalli, P. “Learning Decision Rules from Randomized Iterative Local Search” in Proceedings of the International Conference on Machine Learning, 2002.
55. Seri, S. and Tadepalli, P. “Model-based Hierarchical Average-Reward Reinforcement Learning,” in Proceedings of the International Conference on Machine Learning, 2002.
56. 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.
57. 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.
58. 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.
59. Tadepalli, P. and Dietterich, T. G. “Hierarchical Explanation-Based Reinforcement Learning,” in Proceedings of International Machine Learning Conference, 1997.
60. Liere, R. and Tadepalli, P. “Active Learning with Committees for Text Categorization,” in Proceedings of National Conference on Artificial Intelligence, 1997.
61. Reddy, C., Tadepalli, P. “Learning Goal-Decomposition Rules using Exercises,” in Proceedings of International Conference on Machine Learning, 1997.
62. 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.
63. Ok, D. and Tadepalli, P., “Auto-exploratory Average Reward Reinforcement Learning,” in Proceedings of National Conference on Artificial Intelligence, 1996.
64. Reddy, C., Tadepalli, P. and Roncagliolo, S., “Theory-guided Empirical Speedup Learning of Goal Decomposition Rules,” in Proceedings of International Machine Learning Conference, 1996.
65. Tadepalli, P., “Learning from Queries and Examples with Tree-Structured Bias,” in Proceedings of International Machine Learning Conference, Amherst, MA, 1993.
66. Tadepalli, P., “A Theory of Unsupervised Speedup Learning,” in Proceedings of National Conference on Artificial Intelligence, SanJose, CA, 1992.
67. Tadepalli, P., “A Formalization of Explanation-Based Macro-operator Learning,” in International Joint Conference on Artificial Intelligence, Sydney, Australia, 1991.

68. Tadepalli, P., "Lazy Explanation-Based Learning: A Solution to the Intractable Theory Problem," Proceedings of International Joint Conference on Artificial Intelligence, Detroit, MI, 1989.
69. Natarajan, B. and Tadepalli, P., "Two New Frameworks for Learning," in Proceedings of the International Machine Learning Conference, Ann Arbor, MI, 1988.
70. 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.
71. 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. Hamidi, M., Tadepalli, P., Goetschalckx, R., and Fern, A., "Active Learning of Hierarchical Policies from State-Action Trajectories," in AAAI Workshop in Trajectory-based Behavior Analytics, 2015.
2. Joshi, S., Khardon, R., Tadepalli, P., Fern, A., and Raghavan, A., "Relational Markov Decision Processes: Promise and Prospects," in the workshop on Statistical Relational AI held at National Conference on AI, 2013.
3. Wilson, A., Fern, A., and Tadepalli, P. "A Behavior-Based Kernel for Policy Search via Bayesian Optimization," in the Workshop on Planning and Acting with Uncertain Models at ICML 2011.
4. Wilson, A., Fern, A., and Tadepalli, P. "Transfer Learning in Sequential Decision Problems: A Hierarchical Bayesian Approach," in the Workshop on Unsupervised and Transfer learning at ICML 2011.
5. Natarajan, S., Joshi, S., Tadepalli, P., Kersting, K., and Shavlik, J. W., "Imitation Learning in Relational Domains: A Functional Gradient Boosting Approach," Snowbird Learning Workshop, 2011.
6. Mehta, N., Tadepalli, P., and Fern, A., "Learning Action Models for Planing," Accepted for the Workshop on Learning and Planning at the International Conference on Automated Planning and Scheduling, 2011.
7. Doppa, J. R., Nasresfahani, M., Sorower, M. S., Irvine, J., Dietterich, T. G., Fern, X. and Tadepalli, P. "Learning Rules from Incomplete Examples via Observation Models" accepted for the Workshop on Formalisms and Methodologies in Learning by Reading at IJCAI, 2011.
8. Sorower, M. S., Dietterich, T. G., Doppa, J. R., Tadepalli, P., and Fern, X. "Mention Model for Learning Rules from Incomplete Examples" accepted for the Workshop on Formalisms and Methodologies in Learning by Reading at IJCAI, 2011.
9. Doppa, J. R., Nasresfahani, M., , M. S., Dietterich, T. G., Fern, X. and Tadepalli, P., "Towards Learning Rules from Natural Texts," In Proceedings of Workshop on Formalisms and Methodologies in Learning by Reading at NAACL, 2010.
10. Natarajan, S., Tadepalli, P. Kunapuli, G. and Shavlik, J. "Knowledge Intensive Learning: Directed vs. Undirected SRL Models," in Workshop on Statistical Relational Learning, 2009.
11. Judah, K., Dietterich, T., Fern, A., Irvine, J. Slater, M., Tadepalli, P., Gervasio, M., Ellwood, C., Jarrold, B., Brdiczka, O. and Blythe, J. "User Initiated Learning for Adaptive Interfaces," in the Workshop on Intelligence and Interaction at IJCAI 2009.
12. Mehta, N., Tadepalli, P., and Fern, A., "Learning and Planning with Partial Models," Workshop on Learning Structural Knowledge from Observations at IJCAI 2009.
13. Doppa, J., Yu, J., Tadepalli, P., and Getoor, L., "Chance-Constrained Programs for Link Prediction," in the Workshop on Analyzing Networks and Learning with Graphs at NIPS 2009.

14. 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.
15. 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.
16. 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.
17. 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.
18. 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.
19. Fern, A., Natarajan, S., Judah, K., and Tadepalli, P., "A Decision-Theoretic Model of Assistance," AAAI Workshop on Modeling Others from Observation, 2006.
20. Natarajan, S., Wong, W-K., and Tadepalli, P., "Structure Refinement in First-Order Conditional Influence Language," ICML Workshop on Statistical Relational Learning, 2006.
21. Mehta, N., Natarajan, S., Tadepalli, P., and Fern, A., "Transfer in Hierarchical Variable-Reward Reinforcement Learning," NIPS Workshop on Inductive Transfer, 2005.
22. Mehta, N. and Tadepalli, P., "Multi-agent Shared Hierarchy Reinforcement Learning," Workshop on Rich Representations for Reinforcement Learning, Bonn, Germany, 2005.
23. 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.
24. 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.
25. 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.
26. 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.
27. Liere, R. and Tadepalli, P. "The Use of Active Learning in Text Categorization," In AAAI Spring Symposium on Machine Learning in Information Access, 1996.
28. 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.
29. 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.

30. 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.
31. Tadepalli, P. and Isukapalli, R., "Learning Plan Knowledge Using Simulators," in Proceedings of the workshop on Knowledge Compilation and Speedup Learning, Aberdeen, Scotland, 1992.
32. 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.
33. Tadepalli, P., "Learning with Inscrutable Theories," in International Machine Learning Workshop, Chicago, IL, 1991.
34. Tadepalli, P., "On Quantifying Approximation," Proceedings of the Workshop on Automatic Generation of Abstractions and Approximations at AAAI National Conference, Boston, MA, 1990.
35. Tadepalli, P., "Lazy Explanation-Based Learning," Proceedings of the Workshop on Explanation at AAAI National Conference, Boston, MA, 1990.
36. Tadepalli, P., "Planning in Games Using Approximately Learned Macros," Proceedings of Machine Learning Workshop, Ithaca, NY, 1989.
37. 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. Joshi, S., Khardon, R., Tadepalli, P., Fern, A., Raghavan, A., "Relational Markov Decision Processes: Promise and Prospects," Invited contribution at the workshop on Statistical Relational AI, 2013.
2. 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.
3. 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.
4. Yenong, Q. and Tadepalli, P., "Local search Methods for Job-Shop Scheduling," Oregon State University Technical Report, 94-30-2.
5. Tadepalli, P. and Ok, D., "H-Learning: A Reinforcement Learning Method to Optimize Undiscounted Average Reward," Oregon State University Technical Report, 94-30-1.
6. Dietterich, T., Ok, D., Tadepalli, P. and Zhang, W., "Reinforcement Learning in Scheduling," in Space Operations and Applications Conference, Houston, TX, 1993.
7. Tadepalli, P., "Knowledge-Based Planning in Games," Carnegie Mellon University Technical Report, CMU-CS-89-135, 1989.
8. Tadepalli, P., "Learning Approximate Plans in Games," Ph.D. Thesis, Rutgers University Technical Report, ML-TR-8, 1986.

Students graduated:

1. Rui Qin, M.S., 2016 Project: Extracting Structured Data from Extreme Weather Reports
2. Kranti Kumar Potanapalli, M.S., 2013 (Co-advised with Alan Fern).
Thesis: Coactive Learning for Multi-robot Search and Coverage

3. Janardhan Rao Doppa, Ph.D., 2014 (Co-advised with Alan Fern).
Thesis: Integrating Learning and Search for Structured Prediction (nominated for ACM Distinguished Dissertation Award)
4. Aaron Wilson, Ph.D., 2012 (Co-advised with Alan Fern).
Thesis: Bayesian Methods for Knowledge Transfer and Policy Search in Reinforcement Learning
5. Neville Mehta, Ph.D., 2011.
Thesis: Hierarchical Structure Discovery and Transfer in Sequential Decision Problems
6. Ronny Bjarnason, Ph.D., 2010 (Co-advised with Alan Fern).
Thesis: Monte-Carlo Planning for Probabilistic Domains
7. Scott Proper, Ph.D., 2010.
Thesis: Solving Multiagent Reinforcement Learning.
8. Sriraam Natarajan, Ph.D., 2007.
Thesis: Effective Decision-Theoretic Assistance Through Relational Hierarchical Models.
9. Charles Parker, Ph.D., 2007.
Thesis: Structured Gradient Boosting.
10. Kiran Polavarapu, M.S., 2006.
Project: Event and Sentiment Extraction in the Financial Domain
11. Thierry Donneaugolencer, M.S., 2005
Project: Planning in Multi-agent Partially Observable Domains using Sparse Sampling.
12. Sriraam Natarajan, M.S., 2004.
Thesis: Multi-criteria Average-reward Reinforcement Learning.
13. Nimish Dharawat (M.S.), 2004
Project: Learning Tree Patterns for Information Extraction.
14. Kim Mach, M.S., 2004.
Project: An Experimental Evaluation of Auto-exploratory Average-reward Reinforcement Learning.
15. Hong Tang, M.S., 2002.
Project: Average-reward Reinforcement Learning for Product Delivery by Multiple Vehicles.
16. Sandeep Seri, M.S., 2002.
Thesis: Hierarchical Average-reward Reinforcement Learning.
17. 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.
18. Michael Chisholm, M.S., 1999.
Thesis: Learning Classification Rules by Randomized Iterative Local Search
19. Ray Liere, Ph.D., 1999.
Thesis: Active Learning with Committees: An Approach To Efficient Learning in Text Categorization Using Linear Threshold Algorithms
20. Chandra Reddy, Ph.D., 1998.
Thesis: Learning Hierarchical Decomposition Rules for Planning: An Inductive Logic Programming Approach.
21. Dokyeong Ok, Ph.D., 1996.
Thesis: A Study of Model-based Average Reward Reinforcement Learning.

22. Peter Drake, M.S., 1995.
Thesis: Constructive Induction for Improved Learning of Boolean Functions.
23. Qi Yenong, M.S., 1993.
Project: Local Search Methods for Job-Shop Scheduling.
24. Silvana Roncagliolo, M.S., 1993.
Project: Empirical Speedup Learning of Decomposition Rules for Planning.
25. Ramana Isukapalli, M.S., 1992.
Thesis: Learning Macro-operators for Planning using Simulators.