

## Refereed Articles

- Michail G. Lagoudakis and Ronald Parr, “Least–Squares Policy Iteration,” *Journal of Machine Learning Research (JMLR)*, **4**, 2003, pp. 1107–1149.
  1. Remi Munos, “Error Bounds for Approximate Value Iteration,” Technical Report CMAP 527, *Ecole Polytechnique*, 2004.
- Michail G. Lagoudakis and Ronald Parr, “Reinforcement Learning as Classification: Leveraging Modern Classifiers,” *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, Washington, DC, U.S.A., August 2003, pp. 424–431.
  1. Alan Fern, SungWook Yoon, and Robert Givan, “Approximate Policy Iteration with a Policy Language Bias,” *Proceedings of NIPS\*2003: Neural Information Processing Systems*, Vancouver, BC, Canada, December 2003.
- Michail G. Lagoudakis and Ronald Parr, “Value Function Approximation in Zero–Sum Markov Games,” *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence (UAI-02)*, Edmonton, AB, Canada, August 2002, pp. 283–292.
  1. Nikos Vlassis, *A Concise Introduction to Multiagent Systems and Distributed AI*, Introductory Text, University of Amsterdam, September 2003.
  2. Boddy, M.S., Harp, S.A., and Nelson, K.S., “CLOCKWORK: Requirements Definition and Technology Evaluation for Robust, Compiled Autonomous Spacecraft Executives,” *Final Report, NASA Grant NAG-2-1624*, January 15, 2004.
  3. Erfu Yang and Dongbing Gu “Multiagent Reinforcement Learning for Multi-Robot Systems: A Survey,” Technical Report CSM-404, Department of Computer Science, University of Essex, 2004.
- Carlos Guestrin, Michail G. Lagoudakis, and Ronald Parr, “Coordinated Reinforcement Learning,” *Proceedings of the 19th International Conference on Machine Learning (ICML-02)*, Sydney, Australia, July 2002, pp. 227–234.
  1. Liviu Panait and Sean Luke, “Cooperative Multi-Agent Learning: The State of the Art,” Technical Report GMU-CS-TR-2003-1, Department of Computer Science, George Mason University, 2003.
  2. Stuart Russell and Andrew L Zimdars, “Q-Decomposition for Reinforcement Learning Agents,” *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, Washington, DC, U.S.A., August 2003, pp. 656–663.
  3. Pieter Jan’t Hoen and Sander M. Bohte, “Collective INtelligence with Sequences of Actions,” *Proceedings of the European Conference on Machine Learning*, Dubrovnik, Croatia, September 2003.
  4. Nikos Vlassis, *A Concise Introduction to Multiagent Systems and Distributed AI*, Introductory Text, University of Amsterdam, September 2003.
  5. Srinidhi Varadarajan, Naren Ramakrishnan, and Muthukumar Thirunavukkarasu, “Reinforcing Reachable Routes,” *Computer Networks*, **43**, 2003, pp. 389–416.
  6. Pieter Jan’t Hoen and Sander M. Bohte, “Collective INtelligence with Task Assignment,” *Proceedings of the Annual Machine Learning Conference of Belgium and the Netherlands*, Brussels, Belgium, January 2004.
  7. Jelle R. Kok and Nikos Vlassis, “Sparse Tabular Multiagent Q-learning,” *Proceedings of the Annual Machine Learning Conference of Belgium and the Netherlands*, Brussels, Belgium, January 2004.

8. Jelle R. Kok and Nikos Vlassis, "Sparse cooperative Q-learning," *Proceedings of the 21st International Conference on Machine Learning*, Banff, Canada, July 2004.
  9. Mohammad Ghavamzadeh and Sridhar Mahadevan, "Learning to Act and Communicate in Cooperative Multiagent Systems using Hierarchical Reinforcement Learning," *Proceedings of the Third International Joint Conference on Autonomous Systems and Multi-Agent Systems (AAMAS-2004)*, New York, July 2004.
- Michail G. Lagoudakis and Ronald Parr, "Model-Free Least-Squares Policy Iteration," *Proceedings of NIPS\*2001: Neural Information Processing Systems*, Vancouver, BC, Canada, December 2001, pp. 1547–1554.
    1. Ralf Schoknecht and Artur Merke, "Convergent Combinations of Reinforcement Learning with Linear Function Approximation," *Proceedings of NIPS\*2002: Neural Information Processing Systems*, Vancouver, BC, Canada, December 2002.
    2. Ralf Schoknecht, "Optimality of Reinforcement Learning Algorithms with Linear Function Approximation," *Proceedings of NIPS\*2002: Neural Information Processing Systems*, Vancouver, BC, Canada, December 2002.
    3. Fletcher Lu, "Exploring Model-Based Methods for Active Machine Learning," *Ph.D. Thesis*, School of Computer Science, University of Waterloo, 2003.
    4. Fletcher Lu and Dale Schuurmans, "Model-based least-squares policy evaluation," *Proceedings of the 16th Conference of the Canadian Society for Computational Studies of Intelligence, AI 2003*, Halifax, Canada, June 2003. *Lecture Notes in Computer Science*, Springer-Verlag, Vol. 2671, 2003, pp. 342–352.
    5. Ishai. Menache, Shie Mannor and Nahum Shimkin, "Basis function adaptation in temporal difference reinforcement learning," *Annals of Operations Research, Special Issue on the Cross Entropy method*, 2004.
    6. Hailin Li and Cihan H. Dagli, "Hybrid Least-Squares Methods for Reinforcement Learning," *Proceedings of the 16th International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems, IEA/AIE 2003*, Loughborough, UK, June 2003. *Lecture Notes in Computer Science*, Springer-Verlag, Vol. 2718, 2003, pp. 471–480.
  - Michail G. Lagoudakis and Michael L. Littman, "Learning to Select Branching Rules in the DPLL Procedure for Satisfiability," *Electronic Notes in Discrete Mathematics (ENDM)*, Vol. 9, Elsevier, *LICS 2001 Workshop on Theory and Applications of Satisfiability Testing (SAT-2001)*, Boston, MA, USA, June 2001.
    1. Bryan W. Singer, "Automating the Modeling and Optimization of the Performance of Signal Processing Algorithms," *Ph.D. Thesis*, School of Computer Science, Carnegie Mellon University, December 2001.
    2. Wheeler Ruml, "Adaptive Tree Search," *Ph.D. Thesis*, Division of Engineering and Applied Sciences, Harvard University, May 2002.
    3. Remshagen A, Truemper K, "Learning in a Compiler for MINSAT Algorithms," *Theory and Practice of Logic Programming*, **3**, May 2003, pp. 271–286.
    4. Kevin Leyton-Brown, "Resource Allocation in Competitive Multi-agent Systems," *Ph.D. Thesis*, Department of Computer Science, Stanford University, August 2003.
    5. Amol Dattatraya Mali, "On quantified weighted MAX-SAT," *Decision Support Systems*, 2004.
  - Michail G. Lagoudakis and Michael L. Littman, "Algorithm Selection using Reinforcement Learning," *Proceedings of the 17th International Conference on Machine Learning (ICML-00)*, Stanford, CA, USA, June 2000, pp. 511–518.

1. Bryan Singer and Manuela M. Veloso, "Learning to Generate Fast Signal Processing Implementations," *Proceedings of the 18th International Conference on Machine Learning (ICML-01)*, Williams College, MA, 2001, pp. 529–536.
2. Bryan Singer and Manuela M. Veloso, "Stochastic search for signal processing algorithm optimization," *Proceedings of the 2001 ACM/IEEE conference on Supercomputing (SC-2001)*, Denver, CO, 2001.
3. Bryan W. Singer, "Automating the Modeling and Optimization of the Performance of Signal Processing Algorithms," *Ph.D. Thesis*, School of Computer Science, Carnegie Mellon University, December 2001.
4. Yonghong Peng, Peter A. Flach, Carlos Soares, Pavel Brazdil, "Improved Dataset Characterisation for Meta-Learning," *Proceedings of the 5th International Conference on Discovery Science (DS2002)*, Luebeck, Germany. *Lecture Notes in Computer Science*, Springer-Verlag, Vol. 2534, 2002, pp. 141–152.
5. Yonghong Peng, Peter A. Flach, Pavel Brazdil, Carlos Soares, "Decision Tree-Based Data Characterization for Meta-Learning," in *Proceedings of the ECML/PKDD-02 Workshop on Integration and Collaboration Aspects of Data Mining, Decision Support and Meta-Learning (IDDM-2002)*, Helsinki, Finland, 2002, pp. 111-122.
6. Bohdana Ratitch and Doina Precup, "Characterizing Markov Decision Processes," *Proceedings of the European Conference on Machine Learning*, Helsinki, Finland, August 2002. *Lecture Notes in Computer Science*, Springer-Verlag, Vol. 2430, 2002, pp. 391–404.
7. Bryan Singer and Manuela Veloso, "Automating the Modeling and Optimization of the Performance of Signal Transforms," *IEEE Transactions on Signal Processing*, **50** (8), 2002, pp. 2003–2014.
8. Bryan Singer and Manuela Veloso, "Learning to Construct Fast Signal Processing Implementations," *Journal of Machine Learning Research*, **3** (Dec), 2002, pp. 887–919.
9. Haipeng Guo, "Algorithm Selection for Sorting and Probabilistic Inference: A Machine Learning-Based Approach," *Ph.D. Thesis*, Department of Computing and Information Sciences, Kansas State University, June 2003.
10. Eric Soubeiga, "Development and Application of Hyperheuristics to Personnel Scheduling," *Ph.D. Thesis*, School of Computer Science and Information Technology, University of Nottingham, June 2003.
11. Kevin Leyton-Brown, Eugene Nudelman, Galen Andrew, Jim McFadden, and Yoav Shoham, "A Portfolio Approach to Algorithm Selection," *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, Acapulco, Mexico, August 2003, pp. 1542–1543.
12. Haipeng Guo, "A Bayesian Approach for Automatic Algorithm Selection," *IJCAI03 Workshop on AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Systems*, Acapulco, Mexico, 2003.
13. Kevin Leyton-Brown, "Resource Allocation in Competitive Multi-agent Systems," *Ph.D. Thesis*, Department of Computer Science, Stanford University, August 2003.
14. Kevin Leyton-Brown, Eugene Nudelman, Galen Andrew, Jim McFadden, and Yoav Shoham, "Boosting as a Metaphor for Algorithm Design," *Proceedings of the 9th International Conference on Principles and Practice of Constraint Programming*, Kinsale, County Cork, Ireland, September 2003. *Lecture Notes in Computer Science*, Springer-Verlag, Vol. 2833, 2003, pp. 899–903.
15. Anita Raja, "Meta-Level Control in Multi-Agent Systems," *Ph.D. Thesis*, Department of Computer Science, University of Massachusetts, Amherst, September 2003.
16. Richard Vuduc, James W. Demmel, and Jeff A. Bilmes, "Statistical Models for Empirical Search-Based Performance Tuning," *International Journal of High Performance Computing Applications*, **18** (1), 2004, pp. 65–94.

17. Michael Alan Heffner, "A Runtime Framework for Adaptive Compositional Modeling," *M.Sc. Thesis*, Department of Computer Science, Virginia Polytechnic Institute and State University, May 2004.
  18. Tom Carchrae and J. Christopher Beck "Low Knowledge Algorithm Control," *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI04)*, San Jose, CA 2004.
- Michail G. Lagoudakis and Thomas H. LaBean, "2D DNA Self-Assembly for Satisfiability," *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, Vol. 54, AMS, *Proceedings of the 5th DIMACS Workshop on DNA Based Computers*, MIT, Boston, MA, USA, June 1999, pp. 141–154.
    1. John H. Reif, "Local Parallel Biomolecular Computation," *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, Vol. 48, AMS, 1999, pp. 217–254.
    2. Erik Winfree, Tony Eng, and Grzegorz Rozenberg, "String Tile Models for DNA Computing by Self-Assembly," *Proceedings of the 6th International Workshop on DNA-Based Computers*, Leiden, The Netherlands, June 2000. *Lecture Notes in Computer Science*, Springer-Verlag, Vol. 2054, 2001, pp. 63–88.
    3. John H. Reif, Thomas H. LaBean, and Nadrian C. Seeman, "Challenges and Applications for Self-Assembled DNA Nanostructures," *Proceedings of the 6th International Workshop on DNA-Based Computers*, Leiden, The Netherlands, June 2000. *Lecture Notes in Computer Science*, Springer-Verlag, Vol. 2054, 2001, pp. 173–198.
    4. Paul W.K. Rothmund, "Theory and Experiments in Algorithmic Self-Assembly," *Ph.D. Thesis*, Computer Science Department, University of Southern California, December 2001.
    5. John H. Reif, "The Emerging Discipline of Biomolecular Computation in the US," *Special issue on Biomolecular Computing, New Generation Computing*, edited by Masami Hagiya, Masayuki Yamamura, and Tom Head, Vol. 20, No. 3, 2002, pp. 217–236.
    6. Leonard Adleman, Qi Cheng, Ashish Goel, Ming-Deh Huang, David Kempe, Pablo Moisset de Espanes, and Paul Wilhelm Karl Rothmund, "Combinatorial Optimization Problems in Self-Assembly," *Proceedings of the 34th annual ACM Symposium on Theory of Computing*, Montreal, Quebec, Canada, 2002, pp. 23–32.
    7. John H. Reif, "Molecular Assembly and Computation: From Theory to Experimental Demonstrations," *plenary paper, 29th International Colloquium on Automata, Languages, and Programming (ICALP)*, Malaga, Spain, July 2002. *Lecture Notes in Computer Science*, Springer-Verlag, Vol. 2380, 2002, pp. 1–21.
    8. Leonard Adleman, Jarkko Kari, Lila Kari, and Dustin Reishus, "On the Decidability of Self-Assembly of Infinite Ribbons," *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'02)*, Vancouver, BC, Canada, November 2002, pp. 530–537.
    9. Alessandra Carbone and Nadrian C. Seeman, "Circuits and programmable self-assembling DNA structures," *Proceedings of the National Academy of Sciences*, **99** (20), 2002, pp. 12577–12582.
    10. Takashi Yokomori, "Molecular computing paradigm toward freedom from Turing's charm," *Natural Computing: an international journal*, **1** (4), 2002, pp. 333–390.
    11. Eric Winfree, "DNA Computing by Self-Assembly," *National Academy of Engineering: The Bridge*, **33** (4), 2003, pp. 31–38.
    12. Mark J. Daley and Lila Kari, "DNA Computing: Models and Implementations," *Comments on Theoretical Biology*, **7**, 2002, pp. 177–198.
    13. Erik Winfree and Renat Bekbolatov, "Proofreading Tile Sets: Error Correction for Algorithmic Self-Assembly," *Proceedings of the 9th International Meeting on DNA Based Computers, DNA9*, Madison, WI, USA, June 2003. *Lecture Notes in Computer Science*, Springer-Verlag, Vol. 2943, 2004, pp. 126–144.
    14. Matthew Cook, Paul W.K. Rothmund, Erik Winfree "Self-Assembled Circuit Patterns," *DNA Computing: 9th International Workshop on DNA Based Computers, DNA9*, Madison, WI, USA, June 2003. *Lecture Notes in Computer Science*, Springer-Verlag, Vol. 2943, 2004, pp. 91–107.

15. Gagan Aggarwal, Michael Goldwasser, Ming-Yang Kao, and Robert Schweller, "Complexities for Generalized Models of Self-Assembly," *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA04)*, New Orleans, Louisiana, January 2004.
  16. Ho-Lin Chen, Qi Cheng, Ashish Goel, Ming-Deh Huang, Pablo Moisset de Espanes, "Invadable Self-Assembly: Combining Robustness with Efficiency," *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA04)*, New Orleans, Louisiana, January 2004.
  17. Qi Cheng, Ashish Goel, Pablo Moisset de Espanes, "Optimal Self-Assembly of Counters at Temperature Two," *Proceedings of the 1st Conference on Foundations of Nanoscience: Self-Assembled Architectures and Devices*, April 2004.
  18. Ho-Lin Chen and Ashish Goel, "Error Free Self-Assembly using Error Prone Tiles," *Proceedings of the 10th International Meeting on DNA Computing, DNA10*, University of Milano-Bicocca, June 2004.
- Michail G. Lagoudakis and Anthony S. Maida, "Neural Maps for Mobile Robot Navigation," *Proceedings of the 1999 IEEE International Joint Conference on Neural Networks (IJCNN-99)*, Washington, DC, USA, July 1999.
    1. Frederick L. Crabbe and Michael G. Dyer, "Goal Directed Adaptive Behavior in Second-Order Neural Networks: The MAXSON family of architectures," *Adaptive Behavior*, **8** (2), 2000, pp. 145–168.
    2. Kian Hsiang Low, "Mobile Robots That Learn to Navigate," *Honours Thesis*, Department of Computer Science, School of Computing, National University of Singapore, April 2001.
    3. Chen Weidong, Fan Changhong, and Xi Yugeng, "On-Line Safe Path Planning in Unknown Environments," *Proceedings of the 2003 IEEE International Conference on Robotics and Automation (ICRA-03)*, Taipei, Taiwan, September 2003, pp. 4191–4196.
  - Bill Z. Manaris, Vallane MacGyvers, and Michail G. Lagoudakis, "Universal Access to Mobile Computing Devices through Speech Input," *Proceedings of the 12th International Florida Artificial Intelligence Research Symposium (FLAIRS-99)*, Orlando, FL, USA, May 1999, pp. 286–292.
    1. Christian Buergy, Jan Reinhardt, and James H. Garrett, Jr. "Wearable Computer: Der Entwurf mobiler IT-Systeme zur Unterstuetzung von Arbeitsprozessen im Bauwesen," *Fortschritt-Berichte VDI - Junge Wissenschaftler forschen*, Reihe 4, Nr. 169, Forum Bauinformatik 2001, (Muenchen). VDI Verlag GmbH, Duesseldorf, 2001.

## Other Published Articles

- Michail G. Lagoudakis, Michael L. Littman, and Ronald Parr, "Selecting the Right Algorithm," *Proceedings of the 2001 AAAI Fall Symposium Series: Using Uncertainty within Computation*, Cape Cod, MA, USA, November 2001.
  1. Haipeng Guo, "Algorithm Selection for Sorting and Probabilistic Inference: A Machine Learning-Based Approach," *Ph.D. Thesis*, Department of Computing and Information Sciences, Kansas State University, June 2003.
  2. Haipeng Guo, "A Bayesian Approach for Automatic Algorithm Selection," *IJCAI03 Workshop on AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Systems*, Acapulco, Mexico, 2003.
- Michail G. Lagoudakis and Michael L. Littman, "Reinforcement Learning for Algorithm Selection" (Student Abstract), *Proceedings of the 17th National Conference on Artificial Intelligence (AAAI-00)*, Austin, TX, USA, July 2000, pp. 1081.

1. Anita Raja and Victor Lesser, “Automated Meta-Level Control Reasoning in Complex Agents,” *Proceedings of Eighteenth International Joint Conference on Artificial Intelligence (IJCAI 2003) Workshop on ‘Agents and Automated Reasoning*, Acapulco, Mexico, August 2003.
- Michail G. Lagoudakis and Anthony S. Maida, “Robot Navigation with a Polar Neural Map” (Student Abstract), *Proceedings of the 16th National Conference on Artificial Intelligence (AAAI-99)*, Orlando, FL, USA, July 1999, pp. 965.
    1. Kian Hsiang Low, “Integrated Robot Planning and Control with Extended Kohonen Maps,” *Master’s Thesis*, Department of Computer Science, School of Computing, National University of Singapore, July 2002 (Best M.Sc. Thesis Award 2002-2003).

## Unpublished Articles

- “Model-Free Least-Squares Policy Iteration” (with Ronald Parr), Technical Report CS-2001-05, Department of Computer Science, Duke University, December 2001.
  1. Remi Munos, “Algorithme d’ iteration sur les politiques avec approximation lineaire,” *Journal Electronique d’ Intelligence Artificielle*, 2004.
  2. Remi Munos, “Error Bounds for Approximate Policy Iteration,” *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, Washington, DC, U.S.A., August 2003, pp. 560-567.
  3. Jan Peters, Sethu Vijayakumar, and Stefan Schaal, “Reinforcement Learning for Humanoid Robots,” *Proceedings of the 3rd IEEE-RAS International Conference on Humanoid Robotics*, Karlsruhe, Germany, 2003.
- “Mobile Robot Local Navigation with a Polar Neural Map,” *M.Sc. Thesis*, Center for Advanced Computer Studies, University of Louisiana, Lafayette, 1998.
  1. Omar M. Al-Jarrah, Omar Q. Bani-Melhem, “Building maps for mobile robot navigation using fuzzy classification of ultrasonic range data,” *Journal of Intelligent and Fuzzy Systems*, **11** (3-4), 2001, pp. 171-184.
- “Spatial Knowledge in Humans, Animals and Robots,” 1998.
  1. Andrew J. MacDonald, “Virtual Archaeology: Virtual Reality as a Tool for the Exploration of Architecture,” *M.Sc. Thesis*, Department of Archaeology, University of York, September 2001.
- Michail G. Lagoudakis, “Neural Networks and Optimization Problems - A Case Study: The Minimum Cost Spare Allocation Problem,” Technical Report, The Center for Advanced Computer Studies, University of Southwestern Louisiana, 1997.
  1. Viet-Minh-Nhat Vo, Omar Cherkaoui, “Traffic Switching Optimization in Optical Routing using Hopfield Networks,” *Proceedings of the Second International Conference of French-Speaking Vietnamese Computer Scientists (RIVF’04)*, Hanoi, Vietnam, February 2004.
- “A Hopfield Neural Network for Dynamic Path Planning and Obstacle Avoidance,” 1997.
  1. P. Ritthipravat, T. Maneewarn, D. Laowattana and K. Nakayama, “Obstacle Avoidance Using Modified Hopfield Neural Network for Multiple Robots” *The 2002 International Technical Conference On Circuit/Systems, Computers and Communications*, Phuket, Thailand, Vol. **1**, July 2002, pp. 30-31.
- “Planning and Intelligent Systems: An Introductory Overview,” Technical Report CACS TR-96-2-1, Center for Advanced Computer Studies, University of Louisiana, Lafayette, 1996.

1. Romaric Charton, “Des agents intelligents dans un environnement de communication multimedia : vers la conception de services adaptatifs,” *Ph.D. Thesis*, Departement de Formation Doctorale en Informatique, Universite Henri Poincare - Nancy 1, December 2003.
- “The 0 – 1 Knapsack Problem: A survey,” 1996.
    1. G.W. DePuy, G.E. Whitehouse, and R.J. Moraga, “Using The Meta-Raps Approach to Solve Combinatorial Problems,” *Proceedings of the 2002 Industrial Engineering Research Conference*, May 2002, Orlando, FL.
    2. Peter F. Stadler and Christoph Flamm, “Barrier Trees on Poset-Valued Landscapes,” *Genetic Programming and Evolvable Machines*, 4(1), 2003, pp. 7–20.