Bibliography

Sorin Achim, Peter Stone, and Manuela Veloso. Building a dedicated robotic soccer system. In Proceedings of the IROS-96 Workshop on RoboCup, pages 41–48, Osaka, Japan, November 1996. (pp. 43, 46, 49, 87, 171, 203)

P. Agre and D. Chapman. Pengi: An implementation of a theory of activity. In Proceedings of the National Conference on Artificial Intelligence, pages 268–272, 1987. (p. 20)

Tomohito Andou. Refinement of soccer agents' positions using reinforcement learning. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 373–388. Springer Verlag, Berlin, 1998. (pp. 97, 167, 206)

David Andre and Astro Teller. Evolving team darwin united. In Minoru Asada and Hiroaki Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*. Springer Verlag, Berlin, 1999. (pp. 100, 206)

Elisabeth Andre, Gerd Herzog, and Thomas Rist. On the simultaneous interpretation of real world image sequences and their natural language description: The system soccer. In Proc. of the 8th ECAI, pages 449–454, Munich, 1988. (p. 202)

David Andre, Emiel Corten, Klaus Dorer, Pascal Gugenberger, Marius Joldos, Johan Kummeneje, Paul Arthur Navratil, Itsuki Noda, Patrick Riley, Peter Stone, Romoichi Takahashi, and Travlex Yeap. Soccer server manual, version 4.0. Technical Report RoboCup-1998-001, RoboCup, 1998. At URL http://ci.etl.go.jp/~noda/soccer/server/Documents.html. (pp. 27, 30, 31, 32, 35, 37, 39, 43)

Elisabeth Andre, Gerd Herzog, and Thomas Rist. Generating multimedia presentations for RoboCup soccer games. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 200–215. Springer Verlag, Berlin, 1998. (p. 204)

Sachiyo Arai, Kazuteru Miyazaki, and Shigenobu Kobayashi. Generating cooperative behavior by multi-agent reinforcement learning. In Sixth European Workshop on Learning Robots, Brighton, UK, August 1997. (p. 167)

Ronald C. Arkin. *Towards Cosmopolitan Robots: Intelligent Navigation in Extended Man-Made Environments*. PhD thesis, University of Massachusetts, September 1987. (p. 20)

R.C. Arkin. Cooperation without communication: Multi-agent schema based robot navigation. Journal of Robotic Systems, 9(3):351–364, April 1992. (p. 20)

Ronald C. Arkin. The 1997 AAAI mobile robot competition and exhibition. AI Magazine, 19(3):13–17, 1998. (p. 179)

Neeraj Arora and Sandip Sen. Resolving social dilemmas using genetic algorithms. In Adaptation, Coevolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium, pages 1–5, Menlo Park,CA, March 1996. AAAI Press. AAAI Technical Report SS-96-01. (pp. 188, 191)

Minoru Asada and Hiroaki Kitano, editors. *RoboCup-98: Robot Soccer World Cup II.* Springer Verlag, Berlin, 1999. (p. 174)

Minoru Asada, Shoichi Noda, Sukoya Tawaratsumida, and Koh Hosoda. Purposive behavior acquisition on a real robot by vision-based reinforcement learning. In Proc. of MLC-COLT (Machine Learning Confernce and Computer Learning Theory) Workshop on Robot Learning, pages 1–9, 1994. (pp. 113, 203)

Minoru Asada, Eiji Uchibe, Shoichi Noda, Sukoya Tawaratsumida, and Koh Hosoda. Coordination of multiple behaviors acquired by vision-based reinforcement learning. In Proc. of IEEE/RSJ/GI International Conference on Intelligent Robots and Systems 1994, pages 917–924, 1994. (p. 203)

Minoru Asada, Shoichi Noda, and Koh Hosoda. Action-based sensor space categorization for robot learning. In Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems 1996 (IROS '96), pages 1502–1509, 1996. (p. 203)

Minoru Asada, Yasuo Kuniyoshi, Alexis Drogoul, Hajime Asama, Maja Mataric, Dominique Duhaut, Peter Stone, and Hiroaki Kitano. The RoboCup physical agent challenge: Phase-I. Applied Artificial Intelligence, 12:251–263, 1998. (p. 42)

Tucker Balch and Ronald C. Arkin. Motor schema-based formation control for multiagent robot teams. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 10–16, Menlo Park, California, June 1995. AAAI Press. (pp. 89, 184)

Tucker Balch. *Behavioral Diversity in Learning Robot Teams.* PhD thesis, College of Computing, Georgia Institute of Technology, 1998. (pp. 98, 188, 192, 206)

Mihai Barbuceanu and Mark S. Fox. COOL: A language for describing coordination in multi agent systems. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 17–24, Menlo Park, California, June 1995. AAAI Press. (pp. 194, 195)

R. A. Barman, S. J. Kingdon, J. J. Little, A. K. Mackworth, D. K. Pai, M. Sahota, H. Wilkinson, and Y. Zhang. DYNAMO: Real-time experiments with multiple mobile robots. In Intelligent Vehicles Symposium, pages 261–266, Tokyo, July 1993. (pp. 19, 185, 202)

M. Benda, V. Jagannathan, and R. Dodhiawala. On optimal cooperation of knowledge sources - an empirical investigation. Technical Report BCS–G2010–28, Boeing Advanced Technology Center, Boeing Computing Services, Seattle, Washington, July 1986. (p. 185)

Kim Binsted. Character design for soccer commentary. In Minoru Asada and Hiroaki Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*. Springer Verlag, Berlin, 1999. (p. 204)

Jim Blythe. *Planning Under Uncertainty in Dynamic Domains*. PhD thesis, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, May 1998. (p. 211)

Alan H. Bond and Les Gasser. An analysis of problems and research in DAI. In Alan H. Bond and Les Gasser, editors, *Readings in Distributed Artificial Intelligence*, pages 3–35. Morgan Kaufmann Publishers, San Mateo, CA, 1988. (p. 181)

Mike Bowling, Peter Stone, and Manuela Veloso. Predictive memory for an inaccessible environment. In Proceedings of the IROS-96 Workshop on RoboCup, pages 28–34, Osaka, Japan, November 1996. (p. 68)

J. A. Boyan and M. L. Littman. Packet routing in dynamically changing networks: A reinforcement learning approach. In J. D. Cowan, G. Tesauro, and J. Alspector, editors, *Advances In Neural Information Processing Systems 6*. Morgan Kaufmann Publishers, 1994. (pp. 49, 167)

Justin A. Boyan and Andrew W. Moore. Generalization in reinforcement learning: Safely approximating the value function. In G. Tesauro, D. Touretzky, and T. Leen, editors, *Advances in Neural Information Processing Systems*, volume 7. The MIT Press, 1995. (p. 166)

V. Braitenburg. *Vehicles – experiments in synthetic psychology*. MIT Press, 1984. (p. 224)

Rodney A. Brooks. A robust layered control system for a mobile robot. IEEE Journal of Robotics and Automation, RA-2:14–23, 1986. (pp. 20, 101)

L. Bull, T.C. Fogarty, and M. Snaith. Evolution in multi-agent systems: Evolving communicating classifier systems for gait in a quadrupedal robot. In Stephanie Forrest, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 382–388, San Mateo,CA, July 1995. Morgan Kaufman. (pp. 195, 197)

Hans-Dieter Burkhard, Markus Hannebauer, and Jan Wendler. AT humboldt — development, practice and theory. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 357–372. Springer Verlag, Berlin, 1998. (pp. 112, 204, 220)

Gianni Di Caro and Marco Dorigo. AntNet: Distributed stigmergetic control for communications networks. Journal of Artificial Intelligence Research, 9:317–365, 1998. (p. 167)

Cristiano Castelfranchi. Commitments: From individual intentions to groups and organizations. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 41–48, Menlo Park, California, June 1995. AAAI Press. (pp. 89, 195, 198)

Gordon Cheng and Alexander Zelinsky. Real-time vision processing for a soccer playing mobile robot. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 144–155. Springer Verlag, Berlin, 1998. (p. 204)

Simon Ch'ng and Lin Padgham. From roles to teamwork: a framework and architecture. Applied Artificial Intelligence, 12:211–232, 1998. (p. 205)

Jeffery A. Clouse. Learning from an automated training agent. In Gerhard Weiß and Sandip Sen, editors, *Adaptation and Learning in Multiagent Systems*. Springer Verlag, Berlin, 1996. (pp. 195, 196, 199)

Philip R. Cohen and Hector J. Levesque. Communicative actions for artificial agents. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 65–72, Menlo Park, California, June 1995. AAAI Press. (pp. 195, 196)

Philip R. Cohen, Hector J. Levesque, and Ira Smith. On team formation. In J. Hintikka and R. Tuomela, editors, *Contemporary Action Theory*. Synthese, 1999. to appear. (pp. 20, 53, 55, 88, 89)

R. Conte, C. Castelfranchi, and F. Dignum. Autonomous norm-acceptance. In J. P. Müller, M. P. Singh, and A. S. Rao, editors, *Intelligent Agents V — Proceedings of the Fifth International Workshop on Agent Theories, Architectures, and Languages (ATAL-98)*, volume 1555 of *Lecture Notes in Artificial Intelligence*, pages 99–112. Springer-Verlag, Heidelberg, 1999. (p. 89)

Silvia Coradeschi and Lars Karlsson. A role-based decision-mechanism for teams of reactive and coordinating agents. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 112–122. Springer Verlag, Berlin, 1998. (p. 205)

Silvia Coradeschi and Jacek Malec. How to make a challenging AI course enjoyable using the RoboCup soccer simulation system. In Minoru Asada and Hiroaki Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*. Springer Verlag, Berlin, 1999. (p. 205)

Robert H. Crites and Andrew G. Barto. Improving elevator performance using reinforcement learning. In D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors, *Advances in Neural Processing Systems 8*, Cambridge, MA, 1996. MIT Press. (pp. 166, 167)

Peter Dayan and Geoffrey E. Hinton. Feudal reinforcement learning. In S. J. Hanson, J. D. Cowan, and C. L. Giles, editors, *Advances in Neural Information Processing Systems 5*. Morgan Kaufmann, San Mateo, CA, 1993. (p. 102)

J. Ll. de la Rosa, A. Oller, J. Vehi, and J. Puyol. Soccer team based on agent-oriented programming. Robotics and Autonomous Systems, October 1997. (p. 205)

Keith S. Decker and Victor R. Lesser. Designing a family of coordination algorithms. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 73–80, Menlo Park, California, June 1995. AAAI Press. (pp. 195, 198)

Keith S. Decker. Distributed problem solving: A survey. IEEE Transactions on Systems, Man, and Cybernetics, 17(5):729–740, September 1987. (p. 181)

Keith S. Decker. Distributed artificial intelligence testbeds. In G. M. P. O'Hare and N. R. Jennings, editors, *Foundations of Distributed Artificial Intelligence*, pages 119–138. Wiley Interscience, 1996. (p. 200)

Keith S. Decker. Task environment centered simulation. In M. Prietula, K. Carley, and L. Gasser, editors, *Simulating Organizations: Computational Models of Institutions and Groups*. AAAI Press/MIT Press, 1996. (pp. 56, 88, 212)

Thomas G. Dietterich. The MAXQ method for hierarchical reinforcement learning. In Proceedings of the Fifteenth International Conference on Machine Learning. Morgan Kaufmann, 1998. (p. 102)

Bruce L. Digney. Emergent hierarchical control structures: Learning reactive/hierarchical relationships in reinforcement environments. In Proceedings of the 4th International Conference of Simulation of Adaptive Behavior, pages 363–372. MIT Press, 1996. (p. 102)

Alexis Drogoul and Anne Collinot. Applying an agent-oriented methodology to the design of artificial organizations: A case study in robotic soccer. Autonomous Agents and Multi-Agent Systems, 1(1):113–129, 1998. (p. 205)

Edmund H. Durfee, Victor R. Lesser, and Daniel D. Corkill. Trends in cooperative distributed problem solving. IEEE Transactions on Knowledge and Data Engineering, 1(1):63–83, March 1989. (p. 181)

Edmund H. Durfee. What your computer really needs to know, you learned in kindergarten. In Proceedings of the Tenth National Conference on Artificial Intelligence, Philadelphia, PA, 1992. Morgan Kaufman. Invited Talk. (p. 181)

Edmund H. Durfee. Blissful ignorance: Knowing just enough to coordinate well. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 406–413, Menlo Park, California, June 1995. AAAI Press. (pp. 184, 185)

Maier Fenster, Sarit Kraus, and Jeffrey S. Rosenschein. Coordination without communication: Experimental validation of focal point techniques. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 102–108, Menlo Park, California, June 1995. AAAI Press. (pp. 188, 192)

Tim Finin, Don McKay, Rich Fritzson, and Robin McEntire. KQML: An information and knowledge exchange protocol. In Kazuhiro Fuchi and Toshio Yokoi, editors, *Knowledge Building and Knowledge Sharing*. Ohmsha and IOS Press, 1994. (pp. 62, 89, 194, 195)

Roger Ford, Craig Boutilier, and Keiji Kanazawa. Exploiting natural structure in reinforcement learning: Experience in robot soccer-playing, 1994. Unpublished Manuscript. (p. 202)

Kikuo Fujimura. *Motion Planning in Dynamic Environments*. Springer-Verlag, 1991. (p. 223)

Masahiro Fujita and Koji Kageyama. An open architecture for robot entertainment. In Proceedings of the First International Conference on Autonomous Agents, pages 435–442, Marina del Rey, CA, February 1997. (pp. 56, 90, 204, 212)

Andrew Garland and Richard Alterman. Multiagent learning through collective memory. In Adaptation, Coevolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium, pages 33–38, Menlo Park,CA, March 1996. AAAI Press. AAAI Technical Report SS-96-01. (p. 134)

Erann Gat. Three-layer architectures. In David KortenKamp, R. Peter Bonasso, and Robin Murphy, editors, *Artificial Intelligence and Mobile Robots*, pages 195–210. AAAI Press, Menlo Park, CA, 1998. (p. 101)

M. R. Genesereth and R. E. Fikes. Knowledge interchange format, version 3.0 reference manual. Technical Report Logic-92-1, Computer Science Department, Stanford University, 1992. (pp. 62, 194, 195)

Natalie S. Glance and Tad Hogg. Dilemmas in computational societies. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 117–124, Menlo Park, California, June 1995. AAAI Press. (pp. 188, 191)

Claudia Goldman and Jeffrey Rosenschein. Emergent coordination through the use of cooperative state-changing rules. In Proceedings of the Twelfth National Conference on Artificial Intelligence, pages 408–413, Philadelphia, PA, 1994. Morgan Kaufman. (pp. 184, 186, 188)

Stephen Grand and Dave Cliff. Creatures: Entertainment software agents with artificial life. Autonomous Agents and Multi-Agent Systems, 1(1):39–58, 1998. (p. 195)

John Grefenstette and Robert Daley. Methods for competitive and cooperative co-evolution. In Adaptation, Coevolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium, pages 45–50, Menlo Park,CA, March 1996. AAAI Press. AAAI Technical Report SS-96-01. (pp. 188, 189)

Barbara J. Grosz. Collaborative systems. AI Magazine, 17(2):67–85, Summer 1996. (pp. 20, 53, 88)

J-S. Gutmann, W. Hattzack, I. Herrmann, B. Nebel, F. Rittinger, A. Topor, T. Weigel, and B. Welsch. The CS Freiburg team. In Minoru Asada, editor, *Proceedings of the second RoboCup Workshop*, pages 451–458, Paris, France, July 1998. (pp. 204, 206)

Afsaneh Haddadi. Towards a pragmatic theory of interactions. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 133–139, Menlo Park, California, June 1995. AAAI Press. (pp. 195, 198)

Kwun Han and Manuela Veloso. Physical model based multi-objects tracking and prediction in robosoccer. In Working Note of the AAAI 1997 Fall Symposium. AAAI, MIT Press, 1997. (pp. 46, 48)

Kwun Han and Manuela Veloso. Reactive visual control of multiple non-holonomic robotic agents. In Proceedings of the International Conference on Robotics and Automation, Leuven,Belgium, May 1998. (pp. 204, 224)

Wong-Gie Han, Seung-Min Baek, Tae-Yong Kuc, and Sung Kyun Kwan. Path planning of visual-served multiple mobile robots using the genetic algorithms. In Proceedings of the Micro-Robot World Cup Soccer Tournament, pages 57–63, Taejon, Korea, November 1996. IEEE Robotics and Automation Society. (p. 204)

Thomas Haynes and Sandip Sen. Evolving behavioral strategies in predators and prey. In Gerhard Weiß and Sandip Sen, editors, *Adaptation and Learning in Multiagent Systems*, pages 113–126. Springer Verlag, Berlin, 1996. (pp. 188, 189)

Thomas Haynes, Roger Wainwright, Sandip Sen, and Dale Schoenefeld. Strongly typed genetic programming in evolving cooperation strategies. In Stephanie Forrest, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 271–278, San Mateo, CA, July 1995. Morgan Kaufman. (p. 189)

Kazuo Hirai. Development of the Honda humanoid robot. Presentation at the CMU robotics institute seminar, November 1997. At URL http://www.honda.co.jp/home/hpr/e_news/robot/. (p. 204)

O.E. Holland. Multiagent systems: Lessons from social insects and collective robotics. In Adaptation, Coevolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium, pages 57–62, Menlo Park,CA, March 1996. AAAI Press. AAAI Technical Report SS-96-01. (pp. 184, 186)

Sun-Gi Hong, Tae-Dok Eom, Il-Kwon Jeong, Changkyu Choi, Jin-Ho Shin, and Ju-Jang Lee. Development of soccer-playing robots: Control, cooperation, and strategy. In Proceedings of the Micro-Robot World Cup Soccer Tournament, pages 134–141, Taejon, Korea, November 1996. IEEE Robotics and Automation Society. (p. 203)

Kenji Horikawa, Masaki Aida, and Toshiharu Sugawara. Traffic control scheme under the communication delay of high-speed networks. In Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS-96), pages 111–117, Menlo Park, California, December 1996. AAAI Press. (p. 153)

Tien C. Hsia and Michael Soderstrand. Development of a micro robot system for playing soccer games. In Proceedings of the Micro-Robot World Cup Soccer Tournament, pages 149–152, Taejon, Korea, November 1996. IEEE Robotics and Automation Society. (p. 203)

Marcus J. Huber and Edmund H. Durfee. Deciding when to commit to action during observation-based coordination. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 163–170, Menlo Park, California, June 1995. AAAI Press. (pp. 188, 190)

Bernardo Huberman and Scott H. Clearwater. A multi-agent system for controlling building environments. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 171–176, Menlo Park, California, June 1995. AAAI Press. (pp. 195, 197)

Naomichi Inoue and Samuel Wilkin. Learning system for the formation and the pass-play. Presented at the poster session of the IROS-96 workshop on RoboCup, November 1997. (p. 206)

N. R. Jennings and T. Wittig. Archon: Theory and practice. In N. M. Avouris and L. Gasser, editors, *Distributed Artificial Intelligence: Theory and Praxis*, pages 179–195. Kluwer Academic Press, 1992. (p. 195)

Nicholas R. Jennings, Katia Sycara, and Michael Wooldridge. A roadmap of agent research and development. Autonomous Agents and Multi-Agent Systems, 1(1):7–38, 1998. (p. 181)

Leslie Pack Kaelbling, Anthony R. Cassandra, and Michael L. Littman. Acting optimally in partially observable stochastic domains. In Proceedings of the Twelfth National Conference on Artificial Intelligence, 1994. (p. 166)

Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. Reinforcement learning: A survey. Journal of Artificial Intelligence Research, 4:237–285, May 1996. (pp. 102, 138, 143, 166)

Do-Yoon Kim and Myung Jin Chung. Path planning for multi-mobile robots in the dynamic environment. In Proceedings of the Micro-Robot World Cup Soccer Tournament, pages 127–132, Taejon, Korea, November 1996. IEEE Robotics and Automation Society. (p. 204)

Kyun-Hoon Kim, Kuk-Won Ko, Joo-Gon Kim, Su-Ho Lee, and Hyung Suck Cho. Multiple micro robots playing robot soccer game. In Proceedings of the Micro-Robot World Cup Soccer Tournament, pages 38–43, Taejon, Korea, November 1996. IEEE Robotics and Automation Society. (p. 203)

Jong-Hwan Kim, editor. *Proceedings of the Micro-Robot World Cup Soccer Tournament*, Taejon, Korea, November 1996. (pp. 46, 171)

Hiroaki Kitano, Milind Tambe, Peter Stone, Manuela Veloso, Silvia Coradeschi, Eiichi Osawa, Hitoshi Matsubara, Itsuki Noda, and Minoru Asada. The RoboCup synthetic agent challenge 97. In Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, pages 24–29, San Francisco, CA, 1997. Morgan Kaufmann. (pp. 27, 177)

Hiroaki Kitano, Minoru Asada, Itsuki Noda, and Hitoshi Matsubara. RoboCup: Robot world cup. Crossroads, 4.3, February 1998. (p. 204)

Hiroaki Kitano, editor. *Proceedings of the IROS-96 Workshop on RoboCup*, Osaka, Japan, November 1996. (p. 170)

Hiroaki Kitano, editor. *RoboCup-97: Robot Soccer World Cup I*. Springer Verlag, Berlin, 1998. (p. 172)

John R. Koza. *Genetic Programming.* MIT Press, 1992. (pp. 100, 206)

Michael L. LaBlanc and Richard Henshaw. *The World Encyclopedia of Soccer.* Visible Ink Press, 1994. (pp. 75, 89)

Jean-Claude Latombe. *Robot Motion Planning.* Kluwer, 1991. (pp. 77, 223)

Victor R. Lesser. Multiagent systems: An emerging subdiscipline of AI. ACM Computing Surveys, 27(3):340–342, September 1995. (p. 181)

Victor Lesser. Reflections on the nature of multi-agent coordination and its implications for an agent architecture. Autonomous Agents and Multi-Agent Systems, 1(1):89–112, 1998. (pp. 195, 198)

Ran Levy and Jeffrey S. Rosenschein. A game theoretic approach to the pursuit problem. In Working Papers of the 11th International Workshop on Distributed Artificial Intelligence, pages 195–213, February 1992. (pp. 184, 185)

Long-Ji Lin. *Reinforcement Learning for Robots Using Neural Networks.* PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, 1993. (p. 102)

Michael Littman and Justin Boyan. A distribute reinforcement learning scheme for network routing. Technical Report CMU-CS-93-165, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, July 1993. (pp. 50, 158)

Michael L. Littman. Markov games as a framework for multi-agent reinforcement learning. In Proceedings of the Eleventh International Conference on Machine Learning, pages 157–163, San Mateo, CA, 1994. Morgan Kaufman. (pp. 167, 188, 189, 199, 203, 206)

Sean Luke, Charles Hohn, Jonathan Farris, Gary Jackson, and James Hendler. Co-evolving soccer softbot team coordination with genetic programming. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 398–411, Berlin, 1998. Springer Verlag. (pp. 100, 206)

Andreas Lux and Donald Steiner. Understanding cooperation: an agent's perspective. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 261–268, Menlo Park, California, June 1995. AAAI Press. (pp. 195, 196)

A. K. Mackworth. On seeing robots. In A. Basu and X. Li, editors, *Computer Vision: Systems, Theory, and Applications*, pages 1–13. World Scientific Press, Singapore, 1993. (p. 202)

Pattie Maes and Rodney A. Brooks. Learning to coordinate behaviors. In Proceedings of the Eighth National Conference on Artificial Intelligence, pages 796–802. Morgan Kaufmann, 1990. (p. 102)

Sridhar Mahadevan and Jonathan Connell. Scaling reinforcement learning to robotics by exploiting the subsumption architecture. In Proceedings of the Eighth International Workshop on Machine Learning, pages 328–332, 1991. (p. 102)

Dario Maio and Stefano Rizzi. Unsupervised multi-agent exploration of structured environments. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 269–275, Menlo Park, California, June 1995. AAAI Press. (p. 90)

Maja J. Mataric. Interaction and intelligent behavior. MIT EECS PhD Thesis AITR–1495, MIT AI Lab, August 1994. (pp. 20, 58, 89, 147, 167, 184, 187, 196, 199)

Maja J. Mataric. Learning to behave socially. In Third International Conference on Simulation of Adaptive Behavior, 1994. (pp. 195, 196)

Maja J. Mataric. Designing and understanding adaptive group behavior. Adaptive Behavior, 4(1), December 1995. (p. 101)

Hitoshi Matsubara, Itsuki Noda, and Kazuo Hiraki. Learning of cooperative actions in multi-agent systems: a case study of pass play in soccer. In Adaptation, Coevolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium, pages 63–67, Menlo Park,CA, March 1996. AAAI Press. AAAI Technical Report SS-96-01. (pp. 113, 203)

Hitoshi Matsubara, Ian Frank, Kumiko Tanaka-Ishii, Itsuki Noda, Hideyuki Nakashima, and Koiti Hasida. Automatic soccer commentary and RoboCup. In Minoru Asada and Hiroaki Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*. Springer Verlag, Berlin, 1999. (p. 204)

Akihiro Matsumoto and Hiromasa Nagai. Decision making by the characteristics and the interaction in multi-agent robotics soccer. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 132–143. Springer Verlag, Berlin, 1998. (pp. 205, 206)

Takashi Matsuyama. Cooperative distributed vision. In Proceedings of the First International workshop on Cooperative Distributed Vision, Kyoto,Japan, October 1997. (p. 193)

Andrew Kachites McCallum. Learning to use selective attention and short-term memory in sequential tasks. In From Animals to Animats: Proceedings of the Fourth International Conference of the Simulation of Adaptive Behavior. MIT Press, 1996. (p. 140)

Hironori Mizuno, Masakatu Kourogi, and Yoichi Muraoka. Building shoobot possible to dribble and shoot. In Proceedings of the IROS-96 Workshop on RoboCup, November 1996. (p. 204)

Hironori Mizuno, Masakatu Kourogi, Yukihide Kawamoto, and Yoichi Muraoka. A method applied for soccer's behaviors using proper feedback and feedforward control. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 156–167. Springer Verlag, Berlin, 1998. (p. 204)

Andrew W. Moore and Christopher G. Atkeson. Prioritized sweeping: Reinforcement learning with less data and less time. Machine Learning, 13:103–130, 1993. (pp. 138, 166)

Yishay Mor and Jeffrey Rosenschein. Time and the prisoner's dilemma. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 276–282, Menlo Park, California, June 1995. AAAI Press. (p. 188)

Alexandros Moukas and Pattie Maes. Trafficopter: A distributed collection system for traffic information. At URL http://trafficopter.www.media.mit.edu/projects/trafficopter/, 1997. (p. 193)

Joerg P. Mueller. *The Design of Intelligent Agents: A Layered Approach*. Springer-Verlag, Berlin, 1996. (p. 101)

J. P. Müller. The right agent (architecture) to do the right thing. In J. P. Müller, M. P. Singh, and A. S. Rao, editors, *Intelligent Agents V — Proceedings of the Fifth International Workshop on Agent Theories, Architectures, and Languages (ATAL-98)*, volume 1555 of *Lecture Notes in Artificial Intelligence*, pages 211–225. Springer-Verlag, Heidelberg, 1999. (p. 89)

Rajani Nadella and Sandip Sen. Correlating internal parameters and external performance: learning soccer agents. In Gerhard Weiß, editor, *Distributed Artificial Intelligence Meets Machine Learning*, pages 137–150. Springer-Verlag, 1997. (p. 206)

Hideyuki Nakashima, Itsuki Noda, and Ichiro Ohsawa. Organic programming for multi-agents. In Victor Lesser, editor, *Proceedings of the First International Conference on Multi-Agent Systems*, page 459, San Francisco, CA, 1995. MIT Press. (poster). (p. 205)

Allen Newell and Herbert A. Simon. *Human Problem Solving*. Prentice-Hall, Inc., Englewood Cliffs, NJ, 1972. (pp. 20, 98)

Allen Newell. *Unified Theories of Cognition*. Harvard University Press, Cambridge, MA, 1990. (p. 98)

Itsuki Noda, Hitoshi Matsubara, Kazuo Hiraki, and Ian Frank. Soccer server: A tool for research on multiagent systems. Applied Artificial Intelligence, 12:233–250, 1998. (pp. 25, 27)

Itsuki Noda. Team GAMMA: Agent programming on gaea. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 500–507. Springer Verlag, Berlin, 1998. (p. 205)

Norihiko Ono and Kenji Fukumoto. A modular approach to multi-agent reinforcement learning. In Gerhard Weiß, editor, *Distributed Artificial Intelligence Meets Machine Learning*, pages 25–39. Springer-Verlag, 1997. (p. 167)

E. Pagello, F. Montesello, A. D'Angelo, and C. Ferrari. A reactive architecture for RoboCup competition. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 434–442. Springer Verlag, Berlin, 1998. (pp. 204, 218)

Lynne E. Parker. *Heterogeneous Multi-Robot Cooperation*. PhD thesis, Massachusetts Institute of Technology, 1994. (pp. 20, 90, 101, 195, 198)

H. Van Dyke Parunak, Allen Ward, and John Sauter. A systematic market approach to distributed constraint problems. In Proceedings of the Third International Conference on Multi-Agent Systems, pages 455–456, 1998. (pp. 195, 197)

H. Van Dyke Parunak. Applications of distributed artificial intelligence in industry. In G. M. P. O'Hare and N. R. Jennings, editors, *Foundations of Distributed Artificial Intelligence*, pages 139–164. Wiley Interscience, 1996. (p. 181)

Barney Pell, Douglas E. Bernard, Steven A. Chien, Erann Gat, Nicola Muscettola, P. Pandurang Nayak, Michael D. Wagner, and Brian C. Williams. An autonomous spacecraft agent prototype. Autonomous Robots, 5(1), March 1998. (p. 88)

Yongyuth Permpoontanalarp. Generalised proof-theory for multi-agent autoepistemic reasoning. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 304–311, Menlo Park, California, June 1995. AAAI Press. (pp. 188, 191)

Mitchell A. Potter, Kenneth A. De Jong, and John J. Grefenstette. A coevolutionary approach to learning sequential decision rules. In Stephanie Forrest, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 366–372, San Mateo,CA, July 1995. Morgan Kaufman. (pp. 195, 197)

M. V. Nagendra Prasad, Victor R. Lesser, and Susan E. Lander. Learning organizational roles in a heterogeneous multi-agent system. In Adaptation, Coevolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium, pages 72–77, Menlo Park,CA, March 1996. AAAI Press. AAAI Technical Report SS-96-01. (pp. 188, 192, 199)

Andrew Price, Andrew Jennings, and John Kneen. RoboCup97: An omnidirectional perspective. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 320–332. Springer Verlag, Berlin, 1998. (p. 204)

J. Ross Quinlan. *C4.5: Programs for Machine Learning.* Morgan Kaufmann, San Mateo, CA, 1993. (pp. 93, 95, 116, 119, 122)

Anand S. Rao and Michael P. Georgeff. BDI agents: From theory to practice. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 312–319, Menlo Park, California, June 1995. AAAI Press. (pp. 101, 184, 185, 195, 199)

Matt Ridley. *The Origins of Virtue : Human Instincts and the Evolution of Cooperation.* Viking Press, April 1997.

Jukka Riekki and Juha Roening. Playing soccer by modifying and combining primitive reactions. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 74–87. Springer Verlag, Berlin, 1998. (p. 205)

RoboCup web page, 1997. At URL http://www.robocup.org. (p. 27)

Jeffrey S Rosenschein and Gilad Zlotkin. *Rules of Encounter.* MIT Press, 1994. (pp. 195, 196)

Christopher D. Rosin and Richard K. Belew. Methods for competitive co-evolution: Finding opponents worth beating. In Stephanie Forrest, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 373–380, San Mateo,CA, July 1995. Morgan Kaufman. (pp. 188, 190)

Shounak Roychowdhury, Neeraj Arora, and Sandip Sen. Effects of local information on group behavior. In Adaptation, Coevolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium, pages 78–83, Menlo Park,CA, March 1996. AAAI Press. AAAI Technical Report SS-96-01. (pp. 184, 185)

Michael K. Sahota, Alan K. Mackworth, Rod A. Barman, and Stewart J. Kingdon. Real-time control of soccer-playing robots using off-board vision: the Dynamite testbed. In IEEE International Conference on Systems, Man, and Cybernetics, pages 3690–3663, 1995. (pp. 113, 202)

Michael K. Sahota. Reactive deliberation: An architecture for real-time intelligent control in dynamic environments. In Proceedings of the Twelfth National Conference on Artificial Intelligence, pages 1303–1308, 1994. (pp. 184, 185, 202)

Michael K. Sahota. Dynasim user guide. At URL http://www.cs.ubc.ca/nest/lci/soccer, January 1996. (pp. 113, 203)

Rafal P. Salustowicz, Marco A. Wiering, and Jurgen Schmidhuber. Learning team strategies: Soccer case studies. Machine Learning, 1998. (p. 166)

A. L. Samuel. Some studies in machine learning using the game of checkers. IBM Journal of Research and Development, 3:211–229, 1959. Reprinted in E. A. Feigenbaum and J. Feldman, editors, *Computers and Thought*, McGraw-Hill, New York 1963. (p. 20)

Tuomas W. Sandholm and Robert H. Crites. On multiagent Q-learning in a semi-competitive domain. In Gerhard Weiß and Sandip Sen, editors, *Adaptation and Learning in Multiagent Systems*. Springer Verlag, Berlin, 1996. (pp. 188, 190)

Tuomas Sandholm and Victor Lesser. Coalition formation among bounded rational agents. In Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence, pages 662–669, Los Angeles, CA, 1995. Morgan Kaufman. (pp. 195, 199)

Tuomas Sandholm and Victor Lesser. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 328–335, Menlo Park, California, June 1995. AAAI Press. (pp. 195, 197)

Tuomas Sandholm and Victor Lesser. Advantages of a leveled commitment contracting protocol. In Proceedings of the Thirteenth National Conference on Artificial Intelligence, pages 126–133, Menlo Park,California, 1996. AAAI Press. (pp. 195, 196)

Randy Sargent, Bill Bailey, Carl Witty, and Anne Wright. Dynamic object capture using fast vision tracking. AI Magazine, 18(1):65–72, Spring 1997. (pp. 171, 204)

Paul Scerri. A multi-layered behavior based system for controlling RoboCup agents. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 467–474. Springer Verlag, Berlin, 1998. (p. 205)

Andrea Schaerf, Yoav Shoham, and Moshe Tennenholtz. Adaptive load balancing: A study in multi-agent learning. Journal of Artificial Intelligence Research, 2:475–500, 1995. (pp. 188, 192, 199)

Jurgen Schmidhuber. A general method for multi-agent reinforcement learning in unrestricted environments. In Adaptation, Coevolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium, pages 84–87, Menlo Park,CA, March 1996. AAAI Press. AAAI Technical Report SS-96-01. (pp. 184, 186)

Sandip Sen, editor. *Adaptation, Coevolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium*, Menlo Park,CA, March 1996. AAAI, AAAI Press. AAAI Technical Report SS-96-01. (p. 181)

Robert E. Shapire. The strength of weak learnability. Machine Learning, 5:197–227, 1990. (p. 102)

Onn Shehory and Sarit Kraus. Task allocation via coalition formation among autonomous agents. In Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence, pages 655–661, Los Angeles, CA, 1995. Morgan Kaufman. (pp. 195, 199)

Wei-Min Shen, Jafar Adibi, Rogelio Adobbati, Bonghan Cho, Ali Erdem, Hadi Moradi, Behnam Salemi, and Sheila Tejada. Building integrated mobile robots for soccer competition. In Proceedings of the International Conference on Robotics and Automation, 1998. (p. 204)

Hyun-Sik Shim, Myung-Jin Jung, Heung-Soo Kim, In-Hwan Choi, Woong-Suk Han, and Jong-Hwan Kim. Designing distributed control architecture for cooperative multiagent systems. In Proceedings of the Micro-Robot World Cup Soccer Tournament, pages 19–25, Taejon, Korea, November 1996. IEEE Robotics and Automation Society. (p. 203)

Atsuchi Shinjoh. RoboCup-3D: The construction of intelligent navigation system. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 188–199. Springer Verlag, Berlin, 1998. (p. 204)

Yoav Shoham. Agent-oriented programming. Technical Report CS-1335-90, Computer Science Dept., Stanford University, 1990. (p. 205)

Jaime Simao Sichman and Yves Demazeau. Exploiting social reasoning to deal with agency level inconistency. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 352–359, Menlo Park, California, June 1995. AAAI Press. (pp. 188, 191)

Reid Simmons, Manuela Veloso, and Stephen Smith, editors. *Proceedings of The Fourth International Conference on Artificial Intelligence Planning Systems*, Pittsburgh,PA, June 1998. AAAI Press. (p. 179)

S. P. Singh and R. S. Sutton. Reinforcement learning with replaceing eligibility traces. Machine Learning, 22:123–158, 1996. (p. 166)

Satinder P. Singh, Tommi Jaakkola, and Michael I. Jordan. Learning without state-estimation in partially observable Markovian decision processes. In Proc. 11th International Conference on Machine Learning, pages 284–292. Morgan Kaufmann, 1994. (p. 166)

Reid G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. IEEE Transactions on Computers, C-29(12):1104–1113, December 1980. (pp. 197, 205)

Peter Stone and Manuela Veloso. Beating a defender in robotic soccer: Memory-based learning of a continuous function. In David S. Touretzky, Michael C. Mozer, and Michael E. Hasselmo, editors, *Advances in Neural Information Processing Systems 8*, pages 896–902, Cambridge, MA, 1996. MIT Press. (pp. 113, 203)

Peter Stone and Manuela Veloso. Towards collaborative and adversarial learning: A case study in robotic soccer. In Adaptation, Coevolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium, pages 88–92, Menlo Park,CA, March 1996. AAAI Press. AAAI Technical Report SS-96-01. (p. 203)

Peter Stone and Manuela Veloso. Using machine learning in the soccer server. In Proceedings of the IROS-96 Workshop on RoboCup, pages 19–27, Osaka, Japan, November 1996. (pp. 105, 203)

Peter Stone and Manuela Veloso. Multiagent systems: A survey from a machine learning perspective. Technical Report CMU-CS-97-193, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, December 1997. (p. 181)

Peter Stone and Manuela Veloso. The CMUnited-97 simulator team. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 387–397. Springer Verlag, Berlin, 1998. (p. 83)

Peter Stone and Manuela Veloso. Communication in domains with unreliable, single-channel, low-bandwidth communication. In Alexis Drogoul, Milind Tambe, and Toshio Fukuda, editors, *Collective Robotics*, pages 85–97. Springer Verlag, Berlin, July 1998. (p. 54)

Peter Stone and Manuela Veloso. A layered approach to learning client behaviors in the RoboCup soccer server. Applied Artificial Intelligence, 12:165–188, 1998. (pp. 91, 105, 115)

Peter Stone and Manuela Veloso. Towards collaborative and adversarial learning: A case study in robotic soccer. International Journal of Human-Computer Studies, 48(1):83–104, January 1998. (p. 113)

Peter Stone and Manuela Veloso. Using decision tree confidence factors for multiagent control. In Proceedings of the Second International Conference on Autonomous Agents, 1998. (p. 125)

Peter Stone and Manuela Veloso. Task decomposition and dynamic role assignment for real-time strategic teamwork. In J. P. Müller, M. P. Singh, and A. S. Rao, editors, *Intelligent Agents V — Proceedings of the Fifth International Workshop on Agent Theories, Architectures, and Languages (ATAL-98)*, volume 1555 of *Lecture Notes in Artificial Intelligence*, pages 293–308. Springer-Verlag, Heidelberg, 1999. (p. 53)

Peter Stone and Manuela Veloso. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. Artificial Intelligence, 1999. To appear. (p. 53)

Peter Stone and Manuela Veloso. Team-partitioned, opaque-transition reinforcement learning. In Minoru Asada and Hiroaki Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*. Springer Verlag, Berlin, 1999. Also in *Proceedings of the Third International Conference on Autonomous Agents*,1999. (p. 135)

Peter Stone, Manuela Veloso, and Patrick Riley. The CMUnited-98 champion simulator team. In Minoru Asada and Hiroaki Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*. Springer Verlag, Berlin, 1999. (p. 217)

Peter Stone. Multiagent learning for autonomous spacecraft constellations. In Proceedings of the NASA Workshop on Planning and Scheduling for Space, 1997. (pp. 56, 88, 212)

Devika Subramanian, Peter Druschel, and Johnny Chen. Ants and reinforcement learning: A case study in routing in dynamic networks. In Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, San Francisco, CA, 1997. Morgan Kaufmann. (pp. 52, 167)

Toshiharu Sugawara and Victor Lesser. On-line learning of coordination plans. COINS technical report 93-27, University of Massachussetts Computer Science Department, 1993. (p. 198)

Toshiharu Sugawara and Victor Lesser. Learning coordination plans in distributed OS environments. In Victor Lesser, editor, *Proceedings of the First International Conference on Multi-Agent Systems*, page 462, San Francisco, CA, 1995. MIT Press. (poster). (pp. 195, 198)

R. S. Sutton. Learning to predict by the methods of temporal differences. Machine Learning, 3:9–44, 1988. (pp. 138, 140, 166)

K. Sycara, K. Decker, A. Pannu, M. Williamson, and D. Zeng. Distributed intelligent agents. IEEE Expert, December 1996. (pp. 20, 195, 198)

Katia Sycara. Multiagent systems. AI Magazine, 19(2):79–92, 1998. (p. 181)

Milind Tambe, Jafar Adibi, Yaser Al-Onaizan, Ali Erdem andGal A. Kaminka, Stacy C. Marsela, Ion Muslea, and Marcelo Tallis. Using an explicit model of teamwork in RoboCup-97. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 123–131. Springer Verlag, Berlin, 1998. (p. 205)

Milind Tambe. Recursive agent and agent-group tracking in a real-time , dynamic environment. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 368–375, Menlo Park, California, June 1995. AAAI Press. (pp. 188, 191)

Milind Tambe. Teamwork in real-world, dynamic environments. In Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS-96), Menlo Park, California, December 1996. AAAI Press. (pp. 65, 90)

Milind Tambe. Tracking dynamic team activity. In Proceedings of the Thirteenth National Conference on Artificial Intelligence, Menlo Park, California, 1996. AAAI Press. (pp. 188, 191)

Milind Tambe. Towards flexible teamwork. Journal of Artificial Intelligence Research, 7:81–124, 1997. (pp. 55, 56, 88, 89, 188, 192, 199, 205, 212)

Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In Proceedings of the Tenth International Conference on Machine Learning, pages 330–337, 1993. (pp. 167, 195, 196, 199)

Gerald Tesauro. TD-Gammon, a self-teaching backgammon program, achieves master-level play. Neural Computation, 6(2):215–219, 1994. (pp. 139, 167)

Eiji Uchibe, Minoru Asada, and Koh Hosoda. Behavior coordination for a mobile robot using modular reinforcement learning. In Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems 1996 (IROS '96), pages 1329–1336, 1996. (p. 203)

William T. B. Uther and Manuela M. Veloso. Generalizing adversarial reinforcement learning. In Proceedings of the AAAI Fall Symposium on Model Directed Autonomous Systems, 1997. (p. 206)

Manuela Veloso and Peter Stone. Individual and collaborative behaviors in a team of homogeneous robotic soccer agents. In Proceedings of the Third International Conference on Multi-Agent Systems, pages 309–316, 1998. (pp. 86, 223)

Manuela Veloso and William Uther. The CMTrio-98 Sony legged robot team. In Minoru Asada and Hiroaki Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*. Springer Verlag, Berlin, 1999. To appear. (p. 204)

Manuela Veloso, Peter Stone, and Kwun Han. CMUnited-97: RoboCup-97 small-robot world champion team. AI Magazine, 19(3):61–69, 1998. (pp. 43, 173)

Manuela Veloso, Peter Stone, Kwun Han, and Sorin Achim. The CMUnited-97 small-robot team. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 242–256. Springer Verlag, Berlin, 1998. (p. 43)

Manuela Veloso, William Uther, Masahiro Fujita, Minoru Asada, and Hiroaki Kitano. Playing soccer with legged robots. In Proceedings of IROS-98, Intelligent Robots and Systems Conference, Victoria, Canada, October 1998. (pp. 178, 204)

Manuela Veloso, Michael Bowling, Sorin Achim, Kwun Han, and Peter Stone. The CMUnited-98 champion small robot team. In Minoru Asada and Hiroaki Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*. Springer Verlag, Berlin, 1999. (pp. 44, 110)

Manuela Veloso, Peter Stone, and Michael Bowling. Anticipation: A key for collaboration in a team of agents. In Proceedings of the Third International Conference on Autonomous Agents, 1999. (pp. 77, 87)

Manuela Veloso, Peter Stone, and Kwun Han. The cmunited-97 robotic soccer team: Perception and multi-agent control. Robotics and Automated Systems, 1999. To appear. Also in *Proceedings of the Second International Conference on Autonomous Agents*, May 1998. (p. 43)

Igor Verner. RoboCup: A challenging environment for engineering education. In Minoru Asada and Hiroaki Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*. Springer Verlag, Berlin, 1999. (p. 205)

Adam Walker and Michael Wooldridge. Understanding the emergence of conventions in multi-agent systems. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 384–389, Menlo Park, California, June 1995. AAAI Press. (pp. 188, 192)

Xuemei Wang. Planning while learning operators. In Proceedings of the Third International Conference on AI Planning Systems, May 1996. (pp. 188, 191)

Christopher J. C. H. Watkins. *Learning from Delayed Rewards*. PhD thesis, King's College, Cambridge, UK, 1989. (pp. 93, 138)

Gerhard Weiß and Sandip Sen, editors. *Adaptation and Learning in Multiagent Systems*. Springer Verlag, Berlin, 1996. (p. 181)

Gerhard Weiß. Distributed reinforcement learning. Robotics and Autonomous Systems, 15:135–142, 1995. (pp. 181, 195, 196, 199)

Gerhard Weiß, editor. *Distributed Artificial Intelligence Meets Machine Learning*, volume 1221. Springer Verlag, Berlin, 1997. (p. 181)

David H. Wolpert. Stacked generalization. Neural Networks, 5:241–259, 1992. (p. 102)

Robert S. Woodworth. *Experimental Psychology*. H. Hold and co., New York, 1938. (p. 37)

K. Yokota, K. Ozaki, A. Matsumoto, K. Kawabata, H. Kaetsu, and H. Asama. Omnidirectional autonomous robots cooperating for team play. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 333–347. Springer Verlag, Berlin, 1998. (p. 204)

Dajun Zeng and Katia Sycara. Bayesian learning in negotiation. In Adaptation, Coevolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium, pages 99–104, Menlo Park,CA, March 1996. AAAI Press. AAAI Technical Report SS-96-01. (pp. 195, 197)

Jieyu Zhao and Jurgen Schmidhuber. Incremental self-improvement for life-time multi-agent reinforcement learning. In Proceedings of the 4th International Conference of Simulation of Adaptive Behavior, pages 363–372. MIT Press, 1996. (p. 167)

Gilad Zlotkin and Jeffrey S. Rosenschein. Coalition, cryptography, and stability: Mechanisms for coalition formation in task oriented domains. In Proceedings of the Twelfth National Conference on Artificial Intelligence, pages 432–437, Menlo Park, California, August 1994. AAAI Press. (pp. 195, 199)