

MACHINE LEARNING IN THE FOUR-LEGGED LEAGUE

Stephan K. Chalup and Craig L. Murch

*School of Electrical Engineering & Computer Science, The University
of Newcastle, Callaghan 2308, Australia*

Abstract: The aim of this work is to contribute some insights and a partial overview of how machine learning methods are used in robotics. We first discuss typical general issues in the relationship between robotics and machine learning. Then we focus on projects associated with the RoboCup competition and symposium, and review the extent to which machine learning approaches have been used in the 4-legged league at RoboCup during the years 1998–2003. Further, we summarise the machine learning methods that were used by our own RoboCup team—the NUbots—in 2002/2003.

Keywords: Autonomous mobile robots, learning algorithms, machine learning, optimisation problems, robot control, robot programming, robot vision, robotics, walking.

1. INTRODUCTION

The robot soccer World Cup, RoboCup, first held in 1997, is the premier annual event in adaptive multi-agent systems. “RoboCup is an attempt to promote AI and robotics research by providing a common task for evaluation of various theories, algorithms, and agent architectures” (Kitano *et al.*, 1997). RoboCup’s ultimate goal is, “to build a team of robot soccer players, which can beat a human World Cup champion team (Kitano, 1998)” and “is expected to generate multiple spin-off technologies (Kitano *et al.*, 1997)”. Another long-term vision of many robotics researchers is to have a team of sophisticated autonomous adaptive robots which can explore natural environments and efficiently perform tasks such as search and rescue. This is reflected by the fact that RoboCup currently has two leagues which address search and rescue, one in simulation and one for real world robots. The remaining five leagues of RoboCup (small size, mid size, humanoid, 4-legged, simulation) are robot soccer leagues (see www.robocup.org).

Machine learning research has developed many models, algorithms and techniques which have shown excellent results in areas such as data mining, pattern recognition, and signal processing (Mitchell, 1997). Machine learning methods not only have the potential

to be useful but will, in our opinion, be necessary to solve some of the more challenging robotics tasks, including the above cited long term goals and visions of RoboCup.

Therefore we would like to know to what degree machine learning methods have become accepted by the robotics community. In particular, it would be interesting to evaluate, from the perspective of machine learning, whether the above stated mission of RoboCup has been successful after seven years i.e. whether RoboCup has been able to foster the use of machine learning technology in robotics or whether machine learning research has gained some advancements through RoboCup. It would be interesting to evaluate whether some leagues employ more, or more sophisticated, machine learning techniques than other leagues and whether the different leagues profit from each other. I.e. can learning approaches developed in one league be applied in another league? Does the competitive character of RoboCup eventually inhibit machine learning research (Stone, 2002; Sammut, 2003)? And finally: Have any significant new machine learning techniques been developed in connection with RoboCup?

Associated with RoboCup is a large number of publications in the RoboCup Symposium proceedings

(since 1997), in the individual team reports and in many different journals and conference proceedings related to robotics or machine learning. Given the space constraints of the present conference paper we approach the above questions by giving a partial review focusing only on the 4-legged league since it started in 1998.

The structure of this paper is as follows: We first discuss in Section 2 some general issues in the relationship between robotics and machine learning. In Section 3 we describe the soccer environment and robot platforms of the 4-legged league. Then in Section 4 we review how machine learning was used in the 4-legged league and in Section 5 we report on our own approaches. Finally, in Section 6, we evaluate possible answers to some of the above questions.

2. ROBOT LEARNING

To describe how a robot can acquire skills to perform tasks such as, for example, motor coordination, collision detection, or colour classification we distinguish three possible approaches:

- (1) *Black box approach*: The robot automatically acquires the desired skills from scratch using machine learning methods that are part of its software system. The robot researcher does not need to implement a partial solution on the robot.
- (2) *Grey box approach*: A rough model of the robot and the desired action sequence is given. A machine learning algorithm is employed to fine-tune the parameters of the model and to refine or optimise the robot's behaviour (Millán, 1997).
- (3) *White box approach*: A detailed mathematical model of the robot and its environment is developed. Then the robot is explicitly programmed to perform the desired task. All parameters are "set by hand" i.e. the robotics researcher selects them individually using empirical tests and intuition.

The grey and black box approaches describe what is meant by *robot learning*—an application of machine learning methods to robotics. It is agreed that robot learning presents the biggest and most interesting challenges for machine learning (Connell and Mahadevan, 1993; Mahadevan, 1996). Practical issues that make robot learning so challenging include (Mahadevan, 1996):

- *High noise levels*: Hardware restrictions (e.g. low resolution cameras) often lead to high levels of sensor noise.
- *Unforeseen actions*: Interaction with the real world can demand that the robots cope with situations they were not prepared for. This can lead to unexpected actions.
- *Time and material constraints*: Learning must be achieved in a relatively small number of train-

ing epochs which depend on how fast the real world robot can act.

- *Real-world real-time requirements*: Many real-world situations require that the robot acts/reacts quickly i.e. it must be able to process data in real-time. Suitable learning or adaptive methods must take this into account.
- *Task complexity*: Depending on the complexity of the task (e.g. quadruped walk is extremely complex) simulations or exact (white box) control models may be impossible or inefficient, i.e. on-line training must be conducted.

In the past, the complexity of real world robotics tasks naturally led to the policy: *Always keep system design as clear and simple as possible*. Then, through empirical testing, researchers were able to understand the system and gain experience with its parameters. All parameters could then be chosen by hand (white box approach).

Only recently, the robotics community has become more open to suggestions from the machine learning community to employ learning algorithms (grey box approach) so that robots could be trained on selected aspects of the task and certain parameter sets could be automatically tuned (Kleiner *et al.*, 2002; Riedmiller and Merke, 2002; Demiris and Birk, 2000). This allows for larger parameter spaces and better fine-tuning. Improvements and advantages that robotics research can gain by incorporating suitable machine technology are potentially huge.

If robotics and machine learning are to marry, then as well as the above mentioned challenges for machine learning on the robotics side, there are typical characteristics of the machine learning methodology which are not easy for robotics to deal with:

- *Bias and parameter tuning*: Many of the sophisticated machine learning methods (e.g. for reinforcement learning) are themselves not well-enough understood to be always optimally applicable on a first trial on a real world robotics platform. They often require setting and tuning of critical learning parameters ('magic numbers') and biases without which the algorithms typically will not perform optimally or will not converge in acceptable time. Often only experts with sufficient experience in using a particular type of model or algorithm are able to gain some immediate advantage from its application.
- *Long training times and low convergence rates*: A real world robotics research project typically cannot afford to investigate an algorithm's behaviour in long training and evaluation runs. To be of practical use for robotics the algorithms must come with estimates of convergence rates and training times.
- *Transparency and interpretability of the learning parameters*: Diligent robotics researchers must be very careful about the transparency

and control of the tools and methods employed. Only then can they keep tight control over the behaviour of a complex robotics system and thereby avoid damage to the expensive hardware. For efficient robot learning the researchers must build domain knowledge into a grey box approach and to facilitate this the learning method, its parameters and biases should have an interpretable meaning.

For successful and efficient use of machine learning technology in robotics, future research is therefore needed to address explicitly the *practicability of advanced machine learning algorithms for robotics*.

3. ENVIRONMENT OF THE 4-LEGGED LEAGUE AND THE AIBO ROBOT PLATFORM

The 4-legged league prescribes the use of Sony AIBO robots. The robots are programmed in a C++ software environment to play autonomously robot soccer. No hardware changes are permitted. Information about the initial and the current operating system is available in (Fujita and Kageyama, 1997) and at <http://openr.org>. The AIBOs use a 64-bit RISC processor. The primary sensor is an internal colour camera. Approximate specifications of the different AIBO models used at RoboCup during the years 2000-2004 are displayed in the table below. Note that the details we collected from different team reports and Sony's webpages did not always coincide. Hardware information during the first years of the league was not always available and teams interpreted hardware parameters in different ways. We also have not included specifications for the prototype robots (DRX-720) which were used before 2000.

Model	ERS-110	ERS-210	ERS-210a	ERS-7
CPU	64-bit	64-bit	64-bit	64-bit
clock		192MHz	384MHz	576MHz
memory	8-16MB	32MB	32MB	64MB
camera	CCD	CMOS	CMOS	CMOS
pixels	176×120	176×144	176×144	208×160
frames sec	30	25	25	30
wireless	no	no	812.11b	812.11b
years	2000	2001-2004	2003-2004	2004

The soccer rules in the 4-legged league of RoboCup are only loosely based on real soccer, but the objective of the game is identical. Before 2002, a team consisted of three robots playing on a field of size 180 cm × 280 cm surrounded by white walls. In 2002 the field size was increased to 270 cm × 420 cm and each team could have four robots including the keeper. The green playing surface is carpeted to protect the robots and to allow better grip. The ball is orange. Coloured goals and corner beacons facilitate localisation via the robot's colour camera. More details about the rules and specifications of the environment are available at the RoboCup Legged League web site <http://openr.org/robocup>.

4. REVIEW OF THE USE OF MACHINE LEARNING IN THE 4-LEGGED LEAGUE

The following review includes publications which contain information about the use of machine learning techniques in the Sony 4-legged league of RoboCup 1998-2003. We primarily focus on the RoboCup Symposium proceedings (Asada and Kitano, 1999; Veloso *et al.*, 2000; Stone *et al.*, 2001; Birk *et al.*, 2002; Kaminka *et al.*, 2003; Browning *et al.*, 2004). However, in cases where we are aware of the use of machine learning methods in the 4-legged league and the results have not published in the RoboCup symposium proceeding but the information is available somewhere else, e.g. in the team reports, the team description papers (TDPs) or other publications, then we include them as well.

In 1998, the 4-legged league was an exhibition league composed of three teams. Even at this early stage, machine learning was used in critical areas (Veloso *et al.*, 1998). The team from Carnegie Mellon University used a supervised learning technique involving conjugate gradient descent to determine thresholds for colour classification. Osaka University utilized a behaviour training mechanism: a human controlled the robot playing soccer, while all of the sensory data and the corresponding action performed by the human trainer was recorded. The C4.5 algorithm (Quinlan, 1993) was then used to extract rule sets for performing various actions such as shooting.

1999 saw the 4-legged league become an official RoboCup league. The number of teams increased to nine. The University of Tokyo's team attempted to use a self organising map to enable the robot to "kick the ball where the robot wants to", as well as two neural networks with backpropagation to calculate ideal head pan and tilts based on the position of the ball (Kobayashi and Yuasa, 2000). The University of New South Wales' (UNSW) entry used a two-dimensional polygon growing algorithm to learn how to classify colours in a 2-D colour space (Lawther and Dalglish, 2000). Osaka University developed a technique for using decision/prediction trees to observe landmarks efficiently and make decisions (Mitsunaga and Asada, 2001).

In 2000, the league was expanded to include twelve teams. The UNSW team extended its polygon growing algorithm to account for all three dimensions of the colour space (Hengst *et al.*, 2001b). McGill University used a nearest neighbour interpolation method to assist in colour table generation (Marceau, 2001).

2001 saw the number of teams increase from twelve to sixteen. UNSW switched from a polygon growing algorithm to using the C4.5 algorithm (Quinlan, 1993) to create decision trees able to perform colour classification (Chan *et al.*, 2002). The team from Essex University developed an evolutionary approach to allow their fuzzy logic based behaviour controller

to learn (Gu and Hu, 2002). They also investigated the use of neural networks for colour detection (Hu *et al.*, 2002). Osaka University made use of a genetic algorithm for tuning certain motion parameters (Mitsunaga *et al.*, 2002). Team Cerberus (a joint team from Bulgaria and Turkey) implemented both decision trees and multi-layer perceptrons for colour classification (Akın *et al.*, 2002).

In 2002, the number of teams in the league increased to nineteen. The UNSW team discontinued its use of C4.5 generated decision trees for colour classification and switched to a nearest neighbour learning technique (Wang *et al.*, 2002). The University of Washington team (Crisman *et al.*, 2002), on the other hand, adopted the C4.5 algorithm. The German team experimented with evolutionary algorithms to improve locomotion (Dahm and Ziegler, 2002).

2003 saw the league expand to twenty-four teams. Team Cerberus switched to a C4.5 decision tree approach for colour classification (Akın *et al.*, 2003). The UNSW team also switched back to using the C4.5 algorithm (Chen *et al.*, 2003). The University of Texas at Austin team made use of a nearest neighbour scheme to learn how to classify colours (Stone *et al.*, 2003), while Griffith University made use of covering algorithms to learn decision lists to perform colour classification (Anderson *et al.*, 2003). The team from the University of Chile developed a genetic-based system for the selection and tuning of rules for the detection of the ball, landmarks and goals (Zagal *et al.*, 2004). UNSW began using a multi-dimensional optimisation method to improve its straight-line walking speed (Chen *et al.*, 2003; Kim and Uther, 2003; Sammut, 2003). Following the competition, teams from both the University of Texas and the University of Newcastle developed their own walk optimisation techniques (Kohl and Stone, 2004a; Kohl and Stone, 2004b; Quinlan *et al.*, 2003).

5. HOW THE NUBOTS HAVE EMPLOYED MACHINE LEARNING

The NUBots team collaborates with the University of Newcastle's local machine learning and robotics research group. Machine learning methods were only incorporated into the team's software system if there were strong indications they would have advantages over a direct (white box) approach.

5.1 One-Class Classification with Support Vector Machines

A natural extension of the support vector machine (SVM) algorithm to unlabelled data was proposed by (Schölkopf *et al.*, 2001) who further noted that the new method should have abundant practical applications and could be regarded as an easy-to-use black

box method as soon as questions like the selection of kernel parameters have been solved. As in a standard binary SVM approach, in the one-class SVM approach the data is mapped into a high-dimensional feature space where a separating hyperplane is calculated via a kernel and quadratic programming. The hyperplane is optimised to separate the training data with maximal distance from the origin while the number of outliers is bounded by some $0 < \nu \leq 1$.

(Quinlan *et al.*, 2004; Quinlan *et al.*, 2003) applied this method to the task of colour classification with AIBO robots. An individual one-class SVM was created for each colour label. Scalable, tight fitting cluster boundaries were obtained for each colour cloud in YUV space. The results of this approach were superior to a previous approach using ellipse fitting.

5.2 Evolutionary Hillclimbing for Speed Optimisation

The NUBots' initial hand-tuned parameterised walk engine for the AIBO ERS-210 from 2002 (a loose derivative of (Hengst *et al.*, 2001a)) used inverse kinematics to specify a set of control parameters based on the robot's joint angles. In 2003 this model was refined to include PID values and to allow independent loci for the front and back legs. Each locus is parameterised so that a large variety of suitable shapes is possible.

Since quadruped locomotion has complex dynamics the interpretation and tuning of the 20-100 walk parameters was inefficient by hand. Therefore a modified version of a (1+1)-evolution strategy was applied to optimise the walk parameters for speed (Quinlan *et al.*, 2003). This approach resulted, after a few hours of training, in an increase of about 20% in walking speed over the speed achieved in 2002. To our knowledge walks with speeds up to almost 30 *cm/s* were the fastest walks ever obtained for the AIBO ERS-210a within the 4-legged league (Kohl and Stone, 2004a). The AIBO ERS-7 can walk even faster.

6. DISCUSSION AND CONCLUSION

The review in Section 4 revealed only timid usage of machine learning methods in the now six year old four-legged league. Compared to the total amount of software and methods provided by each team, machine learning played only a minor role. If a team decided to employ a machine learning method it was typically a well-established method such as, for example, the decision tree method C4.5 for which well-tested implementations exist.

This indicates that for the roboticists of the 4-legged league the "practicability of a machine learning method" is an essential condition. "Practicability" means it has been shown that the method works efficiently, a good implementation exists, and it is convenient to use and well-documented. The NUBots

and some other teams investigated the use of several machine learning techniques. A machine learning approach was only selected if it was likely to significantly improve a white box approach, or in situations where the latter was not feasible.

At this point it could seem that RoboCup's attempt to foster research through a soccer competition has not yet been very successful, at least from the perspective of machine learning research. However, RoboCup had seven different leagues in 2003: the two rescue leagues, the 4-legged league, the humanoid league, the soccer simulation league, the F-180 small sized league, and the F-2000 medium sized league.

In the soccer simulation league, matches can be run quickly in the large number necessary for many machine learning algorithms. Therefore machine learning is frequently used and the simulation league has significantly contributed to the development of reinforcement learning in multi-agent domains (Riedmiller and Merke, 2002; Stone and Veloso, 2000).

Among the real robot soccer leagues the F-2000 and F-180 leagues are much more accessible to machine learning approaches than the 4-legged league, with its fully autonomous Sony AIBO robots and their restricted memory and processing power. For example, in the F-2000 league each robot can carry a laptop computer and therefore has enough power to process sophisticated machine learning techniques. This has led to the development of interesting new methods and spin-off projects, see e.g. (Weigel *et al.*, 2002).

We expect that an extension of our partial review, which could also include the other leagues and address RoboCup as a whole, would show that machine learning and robotics profit from each other very well in RoboCup. And the RoboCup competition is a very good benchmark for robotics researchers to compare their own results with the approaches of others. However, it will take a long time before machine learning methods are accepted by roboticists, and until methods which have been successful in one of the other leagues filter through to the 4-legged league. This process could be accelerated if future research explicitly demonstrates how to improve the "practicability of machine learning technology for robotics".

7. ACKNOWLEDGEMENTS

This study was supported by a University of Newcastle RIBG grant, the Centre for Complex Dynamic Systems and Control, the School of Electrical Engineering and Computer Science, the Faculty of Engineering and Built Environment, and a summer scholarship to Craig Murch. We would like to thank Rick Middleton, Michael Quinlan and the NUbots, as well as the other 4-legged teams whose work was the basis of this article.

8. REFERENCES

- Akın, H.L., A. Topalov and O. Kaynak (2002). Cerberus 2001 Team Description. In: Birk *et al.* (2002). pp. 689–692.
- Akın, H.L., M.K. Baloğlu, H.K. Bağcı, S. Bayhan, Ç. Meriçli, D. Poslu, O. Sever, O.T. Yıldız, S. Argirov, B. Marinov, P. Pavlova, N. Shakev, J. Tombakov and A. Topalov (2003). Cerberus 2003 team report. Technical report. Boğaziçi University and Technical University Sofia.
- Anderson, S., K. Croaker, J. Dockter, V. Estivill-Castro, J. Fenwick, N. Lovell and S. Seymon (2003). MiPAL Team Griffith - Summary of our technology and implementation. Technical report. Griffith University.
- Asada, M. and Kitano, H., (Eds.) (1999). *RoboCup-1998: Robot Soccer World Cup II*. Vol. 1604 of LNCS. Springer.
- Birk, A., Coradeschi, S. and Tadokoro, S., (Eds.) (2002). *RoboCup-2001: Robot Soccer World Cup V*. Vol. 2377 of LNCS. Springer.
- Browning, B., Polani, D., Bonarini, A. and Oshida, K., (Eds.) (2004). *RoboCup-2003: Robot Soccer World Cup VII*. LNCS. Springer.
- Chan, S., M. Sio, T. Volgelgesang, T. F. Yik, B. Hengst, S. B. Pham and C. Sammut (2002). The UNSW RoboCup 2001 Sony Legged Robot League Team. In: Birk *et al.* (2002). pp. 39–45.
- Chen, J., E. Chung, R. Edwards, N. Wong, B. Hengst, C. Sammut and W. Uther (2003). Rise of the AIBOS III - AIBO Revolutions. Technical report. University of New South Wales.
- Connell, J. and S. Mahadevan (1993). *Robot Learning*. Kluwer Academic Publishers.
- Crisman, Z., E. Curre, C.T. Kwok, L. Meyers, N. Ratliff, L. Tsybert and D. Fox (2002). Team Description: UW Huskies-02. Technical report. University of Washington.
- Dahm, I. and J. Ziegler (2002). Adaptive methods to improve self-localization in robot soccer. In: Kaminka *et al.* (2003). pp. 393–408.
- Demiris, John and Birk, Andreas, (Eds.) (2000). *Interdisciplinary Approaches to Robot Learning*. Robotics and Intelligent Systems Series. World Scientific.
- Fujita, Masahiro and Koji Kageyama (1997). An open architecture for robot entertainment. In: *Proceedings of the First International Conference on Autonomous Agents*. pp. 435 – 442.
- Gu, D. and H. Hu (2002). Evolving fuzzy logic controllers for Sony legged robots. In: Birk *et al.* (2002). pp. 356–361.
- Hengst, B., D. Ibbotson, S. B. Pham and C. Sammut (2001a). Omnidirectional motion for quadruped robots. In: Birk *et al.* (2002). p. 368.
- Hengst, B., D. Ibbotson, S.B. Pham, J. Dalglish, Mike. Lawther, Phil. Preston and C. Sammut (2001b). The UNSW RoboCup 2000 Sony Legged League Team. In: Stone *et al.* (2001). pp. 64–75.

- Hu, H., D. Gu, D. Golubovic, B. Li and Z. Lio (2002). Essex Rovers 2001 Team Description. In: Birk *et al.* (2002). pp. 697–700.
- Kaminka, Gal A., Lima, Pedro U. and Rojas, Raul, (Eds.) (2003). *RoboCup-2002: Robot Soccer World Cup VI*. Vol. 2752 of LNCS. Springer.
- Kim, M. S. and W. Uther (2003). Automatic gait optimisation for quadruped robots. In: *Australasian Conference on Robotics and Automation (ACRA'2003)*.
- Kitano, H., (Ed.) (1998). *RoboCup-1997: Robot Soccer World Cup I*. Vol. 1395 of LNCS. Springer.
- Kitano, Hiroaki, Minoru Asada, Yasuo Kuyoshi, It-suki Noda, Eiichi Osawa and Hitoshi Matsubara (1997). Robocup: A challenge problem for AI and robotics. In: Kitano (1998).
- Kleiner, A., M. Dietl and B. Nebel (2002). Towards a life-long learning soccer agent. In: Kaminka *et al.* (2003).
- Kobayashi, Y. and H. Yuasa (2000). Team ARAIBO. In: Veloso *et al.* (2000). pp. 758–761.
- Kohl, N. and P. Stone (2004a). Machine learning for fast quadrupedal locomotion. In: *The Nineteenth National Conference on Artificial Intelligence*.
- Kohl, Nate and Peter Stone (2004b). Policy gradient reinforcement learning for fast quadrupedal locomotion. In: *IEEE International Conference on Robotics and Automation*.
- Lawther, M. and J. Dalglish (2000). UNSW United. In: Veloso *et al.* (2000). pp. 788–791.
- Mahadevan, S. (1996). Machine learning for robots: A comparison of different paradigms. In: *Workshop on Towards Real Autonomy, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-96)*.
- Marceau, G. (2001). The McGill's RedDogs Legged League System. In: Stone *et al.* (2001). pp. 627–630.
- Millán, J. del R. (1997). Incremental acquisition of local networks for the control of autonomous robots. In: *7th International Conference on Artificial Neural Networks, Special Session on 'Adaptive Autonomous Agents'*. pp. 739–744.
- Mitchell, Tom (1997). *Machine Learning*. McGraw Hill.
- Mitsunaga, M. and M. Asada (2001). Observation strategy for decision making based on information criterion.. In: Stone *et al.* (2001). pp. 189–198.
- Mitsunaga, N., Y. Nagai, T. Ishida, T. Izumi and M. Asada (2002). BabyTigers 2001: Osaka Legged Robot Team. In: Birk *et al.* (2002). pp. 685–688.
- Quinlan, J. R. (1993). *C4.5: Programs for Machine Learning*. Morgan Kaufmann.
- Quinlan, Michael J., Stephan K. Chalup and Richard H. Middleton (2003). Techniques for improving vision and locomotion on the AIBO robot. In: *Australasian Conference on Robotics and Automation (ACRA'2003)*.
- Quinlan, Michael J., Stephan K. Chalup and Richard H. Middleton (2004). Application of SVMs for colour classification and collision detection with AIBO robots. In: *Advances in Neural Information Processing Systems 16 (NIPS'2003)* (Sebastian Thrun, Lawrence Saul and Bernhard Schölkopf, Eds.). MIT Press. Cambridge, MA.
- Riedmiller, Martin and Artur Merke (2002). Using machine learning techniques in complex multi-agent domains. In: *Perspectives on Adaptivity and Learning* (I. Stamatescu, W. Menzel, M. Richter and U. Ratsch, Eds.). LNCS. Springer.
- Sammut, Claude (2003). Robot soccer: Science or just fun and games?. In: *AI 2003: Advances in Artificial Intelligence 16th Australian Conference on AI* (T. D. Gedeon and L. C. C. Fung, Eds.). Vol. 2903 of LNAI. pp. 12–23.
- Schölkopf, Bernhard, John C. Platt, John Shawe-Taylor, Alex J. Smola and Robert C. Williamson (2001). Estimating the support of a high-dimensional distribution. *Neural Computation* **13**, 1443–1471.
- Stone, P., Balch, T. and Kraetzschmar, G., (Eds.) (2001). *RoboCup-2000: Robot Soccer World Cup IV*. Vol. 2019 of LNCS. Springer.
- Stone, P., K. Dresner, S.T. Erdoğan, P. Fiedelman, N.K. Jong, N. Kohl, G. Kuhlmann, E. Lin, M. Sridharan, D. Stronger and G. Hariharan (2003). UT Austin Villa 2003: A New RoboCup Four-Legged Team. Technical report. University of Texas at Austin.
- Stone, Peter (2002). Multiagent competitions and research: Lessons from RoboCup and TAC. In: Kaminka *et al.* (2003). pp. 224–237.
- Stone, Peter and Manuela Veloso (2000). Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots* **8**(3), 345–383.
- Veloso, M., Pagello, E. and Kitano, H., (Eds.) (2000). *RoboCup-1999: Robot Soccer World Cup III*. Vol. 1856 of LNCS. Springer.
- Veloso, M., W. Uther, M. Fujita, M. Asada and Kitano, H. (1998). Playing soccer with legged robots. In: *Proceedings of Intelligent Robots and Systems Conference (IROS-98)*.
- Wang, Z., J. Wong, T. Tam, B. Leung, M. S. Kim, J. Brooks, A. Chang and N. Von Huben (2002). rUNSWift UNSW RoboCup2002 Sony Legged League Team. Technical report. University of New South Wales.
- Weigel, T., J.-S. Gutman, M. Dietl, A. Kleiner and B. Nebel (2002). CS Freiburg: Coordinating robots for successful soccer playing. *IEEE Transactions on Robotics and Automation* **18**(5), 685–699.
- Zagal, J.C., J. Ruiz-del-Solar, P. Guerrero and R. Palma (2004). Evolving visual object recognition for legged robots. In: Browning *et al.* (2004).