PREPRINT

Machine Learning with AIBO Robots in the Four-Legged League of RoboCup

Stephan K. Chalup, *Member, IEEE,* Craig L. Murch, and Michael J. Quinlan

***Abstract*— Robot learning is a growing area of research at the intersection of robotics and machine learning. The main contributions of this article include a review of how machine learning has been used on Sony AIBO robots and at RoboCup with a focus on the four-legged league during the years 1998–2004. The review shows that the application oriented use of machine learning in the four-legged league was still conservative and restricted to a few well-known and easy to use methods such as standard decision trees, evolutionary hill climbing, and support vector machines. Method oriented spin-off studies emerged more frequently and increasingly addressed new and advanced machine learning techniques. Further the article presents some details about the growing impact of machine learning in the software system developed by authors' robot soccer team—the NUbots.**

***Index Terms*— Machine learning, RoboCup, Sony AIBO, Robot soccer, Robot vision systems, Legged locomotion, Learning systems.**

I. Introduction

THE robot competition and symposium, RoboCup, is the premier annual event in adaptive multi-agent systems. In 1997 it was held for the first time: "RoboCup is an attempt to promote AI and robotics research by providing a common task for evaluation of various theories, algorithms, and agent architectures [61]". Research towards RoboCup's ultimate aim "to build a team of robot soccer players, which can beat a human World Cup champion team [58]"—"is expected to generate multiple spin-off technologies [60]". 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 has two leagues which address search and rescue, one in simulation and one for real world robots. However, the majority of the leagues of RoboCup (simulation, f-2000 middle size, f-180 small size, humanoid, four-legged) are soccer leagues (see `www.robocup.org`).

In the soccer simulation league matches can be run rapidly to generate the large amount of data 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 [4], [98], [109].

Among the real robot soccer leagues, the middle size and the small size [15] leagues are more easily accessible to machine learning approaches than the four-legged league with its fully autonomous Sony AIBO[1] robots and their limited memory and processing power. For example, in the middle size 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 [121].

The four-legged league differs in several respects from the other (soccer) leagues of RoboCup. In contrast to the simulation league and the small size league it is set in a real-world laboratory environment with fully autonomous soccer agents. All teams in the four-legged league must use the same hardware which is the AIBO robot of Sony Corporation [1], [33], [34], [36]. The four-legged league emphasizes the comparison and development of intelligent software systems rather than hardware development although not all teams may use the latest AIBO model. Due to the limited processing power, robot control and team play had to overcome significant challenges during the first years of the four-legged league. However, at the latest competitions the top teams demonstrated exciting game play. The AIBO's smart dog-like design exhibits characteristic artificial creature features which trigger human observers to connect emotionally with the robots. Recent public performances of the "dog teams" attracted cheering crowds of excited robot soccer fans. Currently the robot soccer teams at RoboCup are probably among the most advanced implementations of machine intelligence on robots known to the public.

Machine learning research has developed models, algorithms, and techniques which have shown excellent results and significant improvements in many application areas such as data mining, pattern recognition, signal processing, and robot control, see for example the book by Mitchell [76] or various relevant application papers at the NIPS conferences [84].

The authors are with the School of Electrical Engineering and Computer Science, The University of Newcastle, Callaghan, NSW 2308, Australia, (e-mail: `chalup@cs.newcastle.edu.au`, facsimile: +61 2 492 16929)

[1] AIBO is a trademark of Sony Corporation.

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 and, in particular, for the above cited long term goals and visions of RoboCup. It is therefore relevant to investigate whether RoboCup teams have successfully applied machine learning methods and gained some advantage from incorporating them into their systems. It would also be interesting to know whether machine leaning research has gained some advancements through RoboCup as predicted by the inventors of RoboCup [61], or whether the competitive character of RoboCup eventually inhibits machine learning research in projects associated with the competition [100], [107].

Publications associated with RoboCup have appeared 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. In the present article we extend our pilot study from 2004 [18] and approach the above questions by focusing on the four-legged league since it started in 1998.

There is a similar body of literature covering machine learning in the other leagues of RoboCup. Previous surveys on multiagent systems included links to the history of RoboCup and emphasised involvements of machine learning [4], [110], [111].

The present article is intended for researchers in interdisciplinary fields which combine machine learning and robotics and areas associated with RoboCup [3] or the Federation International Robosoccer Association FIRA [2]. Due to the growing impact of these initiatives the article could also be helpful for researchers with a more general interest in current developments in artificial intelligence or cybernetics.

The structure of this paper is as follows: First, in Section II, some general issues in the relationship between robotics and machine learning are discussed. In Section III the soccer environment and robot platforms are described. Then in Section IV the main machine learning tasks that occur in the four-legged league are explained. The NUbots' approaches are addressed in Section V. The remainder of the article presents a survey describing how machine learning has been used in the four-legged league (Section VI) and on AIBO robots in general (Section VII). Finally, in Section VIII, possible answers to some of the above questions are discussed.

II. Robot Learning

To describe how a robot can acquire skills to perform tasks such as, for example, motor coordination, collision detection, and colour classification, three general approaches can be distinguished. In the *black box* approach a robot automatically and autonomously acquires the desired skills without any prior assumptions on the environment using machine learning methods that are part of its software system. The robot researcher does not need to provide the robot with a partial solution. To indicate the converse of black box we use the term white box. That is, for a *white box* approach a complete mathematical model of the robot and its environment can be developed and no machine learning is required. The robot is explicitly programmed to perform the desired task. All parameters are set by hand, that is, the robotics researcher determines them individually using empirical tests and intuition. In the intermediate situation, the *grey box* approach, a partial model of the environment and the desired action sequence is available. A machine learning algorithm is employed to fine-tune the parameters of the sequence and to refine or optimize the robots behaviour [36]. The grey and black box approaches describe what is meant by *robot learning*—the application of machine learning methods to robotics [22], [29], [32]. In practice at RoboCup all machine learning approaches belong to the grey box category.

There are a number of practical reasons which make the application of machine learning methods to robotics challenging [22], [71], [73]:

- *High noise levels:* Hardware limitations (for example, low camera resolution) often lead to high levels of noise in the data.
- *Stochastic actions:* Interaction with the real world requires that robots cope with situations for which they are not prepared. This can lead to unexpected actions.
- *Time and material constraints:* Learning must be achieved in a relatively small number of training 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. a quadruped walk is extremely complex) simulations or exact (white box) control models are, in many cases, not possible or inefficient, that is, on-line training must be conducted with the real robot.

The complexity of many real-world robotics tasks has naturally led to the use of complicated white box models. Through extensive empirical testing, robotics researchers try to understand the system and gain some insight into the parameters of the algorithms used. Parameters may then be chosen by hand.

Only recently the robotics community has become more open to suggestions from machine learning researchers 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 [28], [62], [98]. This allows for larger parameter spaces and better fine-tuning. The potential improvements and advantages which robotics research can gain by incorporating suitable machine technology are huge.

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

- *Bias and parameter tuning:* Many sophisticated machine learning methods (for example, 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

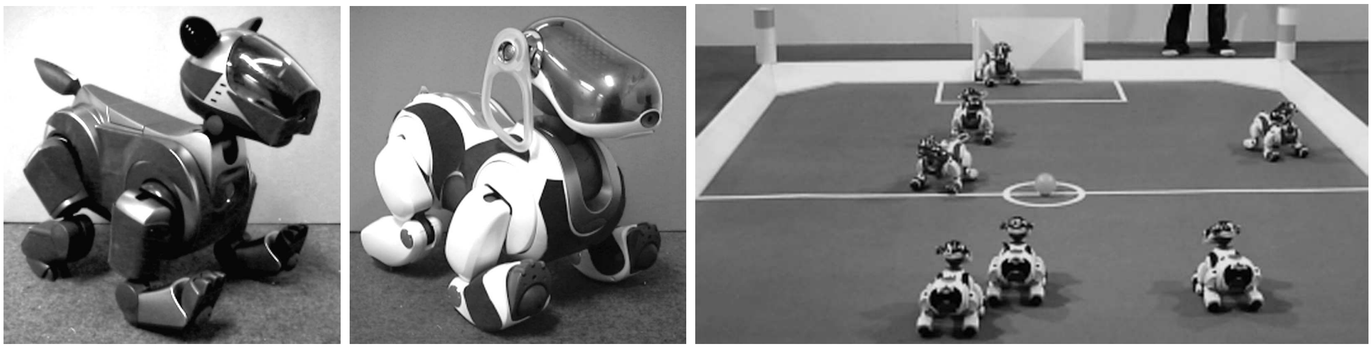

Fig. 1. The dark grey robot on the left is the Sony ERS-210 model first introduced in 2001. The white robot in the middle is the newer Sony ERS-7 robot that has been in use since 2004. Both robots are shown wearing a blue uniform. The right most image shows the robots preparing for a kick-off in a game at RoboCup 2004.

without which the algorithms typically would not perform optimally or would not converge in acceptable time. Often only experts with sufficient experience in using a particular type of model or algorithm have a chance to gain some immediate advantage from its application.

- *Long training times and poor 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 interest for robotics the algorithms must come with practically useful estimates of convergence rates and training times. Research into reducing training times has become a hot topic for real world robotics and recently methods such as model-based reinforcement learning, learning by imitation, or behavioural cloning have been proposed and tested on tasks such as learning helicopter and fixed-wing aircraft control [17], [53], [83].
- *Transparency and interpretability of 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 which is necessary to avoid hardware damage. For efficient robot learning the researchers would have to 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 advised to address explicitly the practicability of machine learning algorithms for robotics.

III. Environment of the Four-Legged League and the AIBO Robot Platform

In the four-legged league, teams must use a hardware platform that is fixed by the rules. This is in contrast to the other robot leagues of RoboCup in which teams construct their own robots. Essentially this means that the four-legged league becomes a software based competition and therefore this league also provides a stable and comparable platform for implementing machine learning on physical robots. The currently allowed hardware consists of several different models of Sony AIBO robots: the ERS-210, ERS-210A and the newer ERS-7 models. All allowed models have 64-bit MIPS processors with clockspeeds of 192MHz (ERS-210), 384MHz (ERS-210A) and 576MHz (ERS-7). Note that the ERS-210 and ERS-210A are identical apart from their processors, so we will only describe the physical specifications of the latter. The robots are programmed in `C++` using Sony's OPEN-R software development kit (see the OPEN-R web-site: `openr.aibo.com`) [35], [59]. The ERS-210A measures (width × height × length) $154mm \times 266mm \times 274mm$, while the ERS-7 is slightly larger at $180mm \times 278mm \times 310mm$. The ERS-210A weighs approximately $1.4kg$, while the ERS-7 is $1.7kg$. The robots are autonomous, but can communicate through wireless LAN (IEEE 802.11b) with the other robots on their team. Wireless communication was introduced into the competition in 2002.

Approximate specifications of the different AIBO models used at RoboCup during the years 2000-2004 are summarised in Table I. Please note that the details collected from different team reports, Sony's webpages, and our own experience, did not always coincide. Information regarding the robot hardware used in the first few years of the league is not readily available, and teams seem to interpret hardware parameters in different ways. We also have not included specifications about any prototype robots (for example the DRX-720 or MUTANT) which were used before 2000 [33], [34], [36]. The soccer rules in

TABLE I
SPECIFICATIONS OF THE DIFFERENT MODELS OF THE SONY AIBO ROBOT AS USED IN THE FOUR-LEGGED LEAGUE OF ROBOCUP

Model	ERS-110	ERS-210	ERS-210A	ERS-7
clock	-	192MHz	384MHz	576MHz
memory	8-16MB	32MB	32MB	64MB
camera	CCD	CMOS	CMOS	CMOS
pixels	176×120	176×144	176×144	208×160
$\frac{\text{frames}}{\text{second}}$	30	25	25	30
years	2000	2001-2004	2003-2004	2004

the four-legged league of RoboCup are only loosely based on real soccer, but the objective of the game is identical. Before

Fig. 2. On the left is the original, unprocessed image obtained from the robot's camera. The middle figure is the same image after colour classification has been performed. The figure on the right shows coloured blobs that have been formed based on the colour classified image. (For printing this article colours have been converted into greyscale.)

2002, a team consisted of three robots playing on a field of size $180cm \times 280cm$ surrounded by white walls. In 2002, the field size was increased to $270cm \times 420cm$ and each team could now have four robots including the goalkeeper. The green playing surface itself 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 detailed rules and specifications of the environment are available at the RoboCup Legged League web site `www.tzi.de/4legged/`. For RoboCup 2005 the field was further enlarged to $4m \times 6m$, the field boundaries were removed, and several other smaller rule changes were introduced.

IV. Principal Machine Tasks in the Four-Legged League

Machine learning methods can potentially be applied to many tasks in the four-legged league, including vision, localisation, locomotion, and behaviour. So far, however, machine learning has been used primarily for colour classification in vision and to improve locomotion.

A. Colour Classification

Robot vision systems are often required to identify landmarks relevant to the operation of the robot. In some cases, colour alone can be used to identify landmarks. For other objects and landmarks, edge detection and shape recognition techniques can be used.

Currently, colour is the primary criterion used to identify landmarks and objects on the four-legged league soccer field. Colour classification is therefore a critically important part of the vision system. Generally speaking, colour classification on AIBO robots [16], [92] is performed by using a pre-computed table (a *colour table*) which maps raw colour information from the YUV colour space into a small set of colours of interest to the robot. Colours of interest to the robot are often termed *classified colours*. Typically, clumps of classified colour in the image are formed into 'blobs'. These steps are illustrated in Figure 2. Blobs are then processed by various *adhoc* techniques to determine which objects appear in the image.

Since the robot is extremely reliant on colour for object detection, a new colour table has to be generated after any change in lighting conditions. This is usually a manual task which requires a human to capture hundreds of images and assign a colour label on a pixel-by-pixel basis. It takes several hours to construct a new colour table using this manual method, yet the table will still contain holes—unmarked points which are surrounded (or almost surrounded) by marked points of a particular colour—and classification errors. Several machine learning algorithms are currently in use by RoboCup teams in order to reduce classification errors and speed up the process. We address these further in section VI.

B. Walk Optimisation

The majority of teams in the legged league use an omnidirectional parameterised walk inspired by the work of Hengst *et al.* [46]. This original walk was commonly known as "PWalk" and some of its characteristics had been adopted by other teams when they developed their own walk engine. In all these walks the end of each paw is commanded to follow a trajectory with inverse kinematics used to calculate the joint angles required to achieve the positions.

The optimisation of the vector of walk parameters has become one of the primary applications areas of machine learning in the four-legged league. Table II gives descriptions of 17 parameters similar as they were used in PWalk. Eight of these parameters affect the stance of the robot (front and back are shown together), five parameters control the type, size, direction, and speed of each step and four parameters control the movement and position of the head.

C. Learning of Team Behaviour

In contrast to some of the other leagues, for example the simulation league, there is a lack of machine learning applied to team behaviour in the four-legged league. This is surprising, since many researchers in this league come from a machine learning background. One possible reason for this is that, against common sense (soccer expert) intuition, team behaviour was rarely of critical importance in match play: low level skills were still the main differentiating factor between most teams in the league.

TABLE II
DESCRIPTION OF PWALK PARAMETERS (LOOSELY BASED ON [46])

hF/hB	Height of front/back hip above the ground (mm)
fsO/bsO	Sideways distance between shoulder and paw (mm)
ffO/bfO	Distance of the paw from shoulder to front (mm)
fdH/bdH	Max height the paw will be lifted off the ground (mm)
walkType	Defines the type or walk action to be used (int)
PG	Time taken for a step in (0.008 second units)
Forward	Distance the robot should walk forwards in step (cm)
Turn	Angle (positive) the robot should turn (degrees)
Left	Distance to move left during the step (cm)
hType	Type of head movement: none, relative, absolute (int)
tilt	Effects the up/down movement of the head (degrees)
pan	Effects the left/right movement of the head (degrees)
mouth	Angle to open the mouth (degrees)

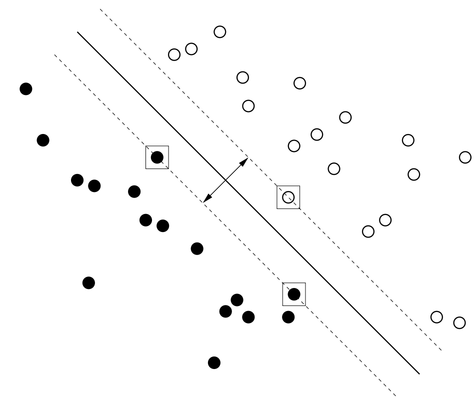

Fig. 3. Maximum Margin Classifier: Separating hyperplane where the emphasised inputs are support vectors.

It may be feasible for four-legged league teams to employ algorithms from other leagues (particularly simulation [109] and small size [14], [15]), although it is somewhat questionable as to whether these methods will transfer to the four-legged environment. In particular, the complexity of the construction of the four-legged league robots makes a sufficiently accurate simulation of the robots and their interaction with the environment difficult.

D. Localisation and World Modeling

In the four-legged league the primary sensor information comes from colour vision which is very noisy. For localisation the robots can use estimates of the location of the coloured corner beacons, the coloured goals, and eventually the white field lines. Kalman filters and particle filters are so far the dominating techniques for localisation and are used by most teams of the four-legged league; see, for example [66], [74]. However, these statistical methods typically are not counted as core machine learning methods.

V. HOW THE NUBOTS HAVE EMPLOYED MACHINE LEARNING

Machine learning methods were only incorporated into the NUbots' software system if there were strong indications it would have advantages over a direct (white box) approach. We describe two examples.

A. One-Class Classification with Support Vector Machines

Support vector machines (SVM) emerged from the field of statistical learning theory [13], [23], [89], [115], [116]. They are now commonly employed for tasks such as classification of handwriting and faces [101]. In their simplest version SVMs can be described as maximum margin classifiers for binary classification (Figure 3). This version only works on linearly separable data by selecting the hyperplane that separates the two classes by the largest margin [25]. A natural extension of the support vector machine algorithm to unlabelled data was proposed by Schölkopf *et al.* [102] who further noted 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. In the one-class SVM approach the data is implicitly 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 parameter $0 < \nu \leq 1$.

Michael Quinlan *et al.* [96], [97] applied support vector machines 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 (see Figure 4). The results of this approach were superior to a previous approach using ellipse fitting [97]. The technique was also applied to the task of collision detection [95] where the one-class SVM is employed as a novelty detection mechanism. In this implementation each training point is a vector containing thirteen elements. These include five walk parameters, along with a sensor reading from the abductor and rotator joints on each of the four legs. The training set is generated by having the robot behave normally (takes approximately 10 minutes) on the field but with the stipulation that all collisions are avoided. Upon training the SVMs decision function will return +1 for all values that relate to a "normal" step, and -1 for all steps that contain a fault. The trained classifier analyses the on-line stream of joint data measurements in samples of ten consecutive data points. If more than two points in one sample are classified as -1 a collision is declared to be detected.

B. Evolutionary Hillclimbing for Speed Optimisation

The robot's legs each follow a trajectory (or locus) in 3-dimensional coordinate space. The model includes PID values and allows independent loci for the front and back legs. Each locus is parameterised so that a large variety of suitable shapes are possible. Since quadruped locomotion has complex dynamics the interpretation and tuning of the 20-100 walk parameters was impossible by hand and a modified version of a (1+1)-evolution strategy was applied to optimise the walk parameters for speed [96].

For our walk engine the parameters to be tuned are defined by a vector θ consisting of 11 walk parameters (turn and strafe are excluded from the learning) and the critical points defining

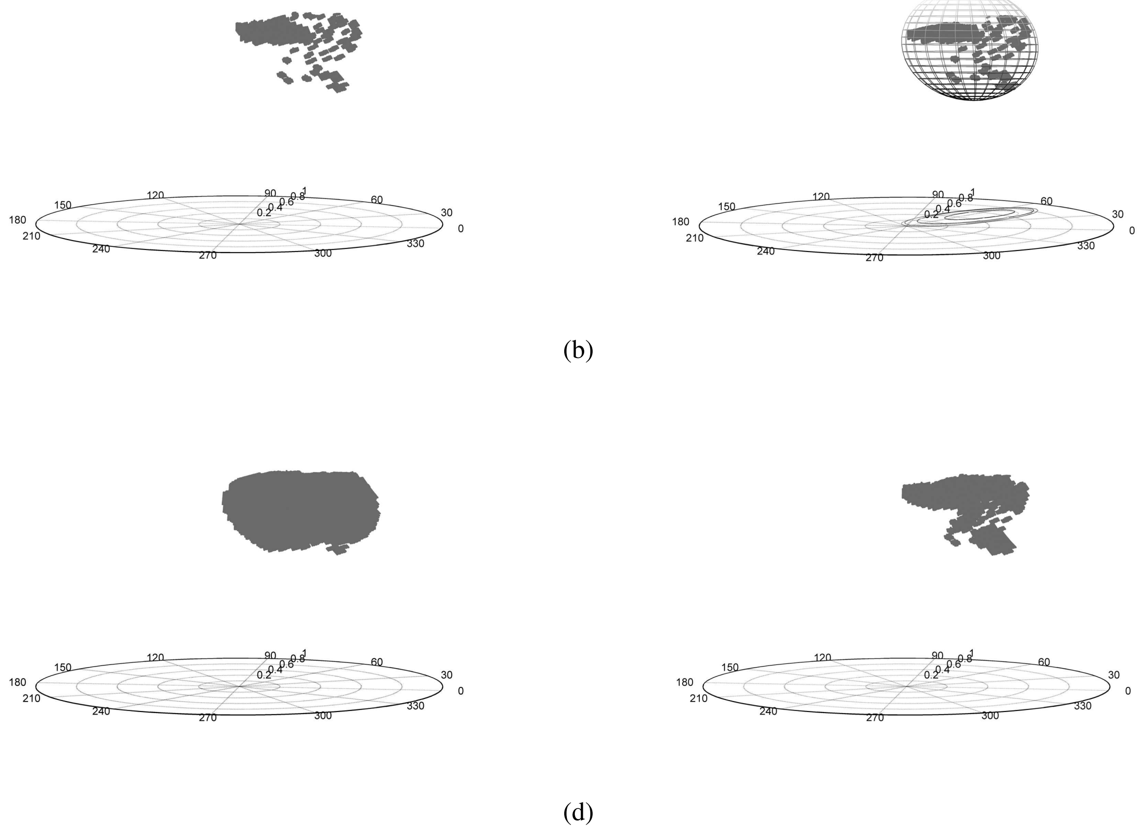

Fig. 4. Colour Classification: (a) Points manually classified at white. (b) Ellipsoid fitted to these white points. (c) Loose fit with one-class SVM technique, ν=0.025 and γ=10. (d) Tight fit with one-class SVM technique, ν=0.025 and γ=250.

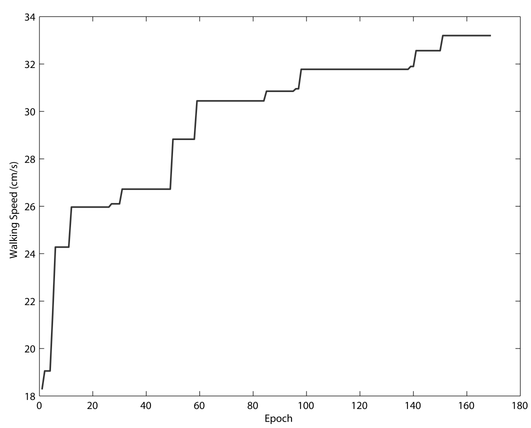

Fig. 5. Example of speed improvement obtained during approximately one hour of training.

an arbitrary locus shape (40 parameters). Each parameter is randomly set to an initial value; for our task we make sure this initial vector is feasible (i.e. it is one that will cause the robot to move in the required direction).

After a few hours of training the learning approach resulted in an about 20% increase in walking speed over the speed achieved in 2002. At that particular time walks with speeds up to almost $30cm/s$ were the fastest walks ever obtained for the AIBO ERS-210(A) within the four-legged league [64]. Upon receiving the ERS-7 we ran the learning algorithm on the new robot from a set of parameters that were developed on the ERS-210 robot. We managed to learn walks with an approximate speed of $41cm/s$ (the initial speed was below $25cm/s$). An example learning curve of this speed optimisation process can be seen in Figure 5.

VI. Machine Learning in the Four-Legged League

The following review includes publications which contain information about the use of machine learning techniques in the Sony four-legged league of RoboCup 1998-2004. We primarily focus on the RoboCup Symposium proceedings [8], [12], [54], [82], [90], [105], [117]. However, in cases when we are aware of the use of machine learning methods in the four-legged league and the results have been published, for example, in the team reports, the team description papers (TDPs), or other publications, then we included them as well. As a symptom of the interdisciplinary character of RoboCup many of the relevant publications are not specialised on machine learning. They often only briefly indicate when a machine learning method was applied without giving details of the algorithms, performance results, or comparisons with alternative approaches. Consequently large parts of the following survey had to remain restricted to showing chronologically how the application of machine learning developed within the four legged league from 1998 to 2004 and only for selected publications some details and discussion of the approach could be included.

A. Paris 1998

In 1998, the four-legged league was an exhibition league by Sony and was composed of three teams from Carnegie Mellon University (CMU), Osaka, and Paris. At this early stage expectations into robot learning were high and machine learning approaches were implemented for critical areas such as colour classification, localisation, and behaviour acquisition [118].

The team from CMU used supervised learning on twenty training images based on a conjugate gradient descent technique to determine thresholds for classification in YUV colour space [119]. For localisation they employed a Bayesian probabilistic approach [118].

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 [94] was then used to extract rule sets for performing various actions such as shooting.

However, it soon became clear that these tasks were extremely challenging and more research into robot learning would be required before satisfactory results could be expected.

B. Stockholm 1999

1999 saw the four-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 "kick the ball where the robot wants to". They also investigated how to train two neural networks with backpropagation to calculate ideal head pan and tilts depending on the position of the ball [63].

Instead of performing the complex calculations of a global self-localisation method the team from Osaka University [80] developed a technique for using decision trees and prediction trees to observe the beacons, the goals, and the ball efficiently and then, without calculating explicit positions, to decide directly which actions to take. They restricted the action space to a small number of elementary actions. The decision and prediction trees of their probabilistic action selection used the ID3 [93] information criterion [77].

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 [27], [68].

C. Melbourne 2000

In Melbourne, the four legged league expanded to include twelve teams.

The team from UNSW showed superior performance in all its matches. Its success was partially due to a new and fast localisation method which employed a stochastic gradient descent learning algorithm for incremental position updates [88]. Given the robot's position and heading in the current world model and a newly perceived position the algorithm updated the current position and heading variables by moving them in small steps in direction towards a newly perceived position. For the method to calculate the perceived position and to perform its updates it was sufficient if only one landmark at a time was visible. This is important because the AIBOs have a very narrow field of view. The update rate can be as high as the frame rate. With each frame the learning algorithm's step size was varied depending on landmark associated decaying confidence parameters [27]. For colour classification the UNSW team extended their polygon growing algorithm from previous year to account for all three dimensions of the colour space [45].

The team from McGill University employed a nearest neighbour interpolation method in YUV colour space to assist in colour table generation for colour classification [72]. This approach allowed them to limit the necessary sample size for "colour training" to 30 pictures which reduced the time required for vision calibration.

D. Seattle 2001

2001 saw the number of four-legged league teams increase from twelve to sixteen.

The Essex Rovers team from the UK developed an evolutionary approach to allow their fuzzy logic based behaviour controller to learn [39]. They also investigated [49]–[51] the use of neural networks for colour detection tables. Their aim was to adapt to changing lighting conditions through variable threshold prediction in YUV space. Their approach was based on a method previously presented by the UNSW team.

In 2001 the UNSW team changed its method for colour classification [19] and switched from the previously used polygon growing algorithm to using the C4.5 decison tree algorithm [94].

The team of Osaka University mentions [79] that they made use of a genetic algorithm for tuning certain motion parameters.

Team Cerberus (a joint team from Bulgaria and Turkey) implemented both decision trees and multi-layer perceptrons for colour classification as well as a genetically trained fuzzy-neural network for behaviour control [6].

E. Fukuoka 2002

At Fukuoka in Japan the number of teams in the four-legged league increased to nineteen.

The UNSW team discontinued their use of C4.5 generated decision trees for colour classification and switched to a nearest neighbour learning technique [120].

The University of Washington team [24], on the other hand, mentioned they had adopted decision trees for colour cluster generation in YUV space. Their main focus, however, was on state estimation and world modelling where they used particle filters to estimate the robots' positions and Kalman filters to trace the ball [42]. They were able to track the position of the robot using only 30 samples with an accuracy of the position estimates of about $10cm$ on average. After the robot was picked up and placed back onto the field at a new location it took them less than 2 seconds to re-localize the robot [24].

The Essex team continued its earlier experiments on behaviour learning with fuzzy logic controllers which allowed the robot to use information from the camera to approach the ball and shoot at the goal. In this study [40] Gu and Hu employed a combination of reinforcement learning [10] and evolutionary computation to learn refinements of the initially handcrafted parameters of their fuzzy logic controller architecture.

Dahm and Ziegler [26] from the German Team envisioned how evolutionary computation could improve localisation. Genetic programming allowed them to increase the robots' walking speed to over $20cm/s$. They proposed that robustness and reliability of the walk could be improved using a multi-objective fitness function. The resulting more reliable odometry together with improved vision and inter robot communication would lead to increased accuracy within their Bayesian probabilistic approach to localisation.

F. Padova 2003

At RoboCup 2003 in Italy the league expanded to twenty-four teams.

For the task of colour classification a variety of machine learning techniques were used by the different teams. In 2003 Team Cerberus [5] and the UNSW team [20] both decided to employ standard decision tree approaches using the C4.5 algorithm [94]. The University of Texas at Austin team made use of a basic nearest neighbour scheme to learn how to classify colours [106]. The team from Griffith University implemented covering algorithms to learn decision lists to perform colour classification [7]. The NUbots demonstrated for the first time that support vector machines [13], [89], [101], [115] can be implemented on AIBOs to achieve refined clustering results in colour space [96], [97]. Although the teams employed a variety of different methods and accordingly obtained classification results of different quality it became clear that sophisticated object recognition with the AIBO camera not only depended on good colour classification but other components of the vision system needed to be improved as well.

Researchers associated with the team from Essex, UK, proposed an adaptive colour segmentation algorithm with the aim to be able to adjust to different lighting conditions. In their pilot experiments [70] they employed self-organising maps for colour segmentation and similar as in previous years an artificial neural net with supervised learning for adaptive threshold selection.

The team from the University of Chile (UChile) applied machine learning to a slightly different but important aspect of the object recognition task. They developed a system for automated selection and tuning of rules for the detection of the ball, the landmarks, and the goals [125]. The system employed a genetic algorithm, a specifically adapted fitness function, and supervised learning on a set of about 180 pre-classified images to evolve successful recognition rules.

The rUNSWift team began using a multi-dimensional optimisation method to improve their straight-line walking speed [20], [57], [100]. Their walk engine implemented, similar as in the previous year [46], a trot gait where diagonally opposite legs were synchronised and the two pairs of synchronised legs were 180^o out of phase. Each of the robot's feet followed a trapezoidal trajectory which was defined by its four corner points each of which could be displaced in three dimensions. Two independent trajectories were used, one for the front legs and one for the back legs. With these constraints the problem could be represented as minimisation over a 24 dimensional search space. Starting from a well selected set of search directions [57] the rUNSWift team employed Powell's multi-dimensional directional set optimisation method [91] which searches along each of these 24 directions separately. UNSW's researchers demonstrated on the soccer fields in Padova how to evaluate the different walk parameter settings by measuring the time required by the physical AIBO robot to run across the field between two opposite beacons at the sidelines. Although the speed achieved by UNSW's initial experiments was limited to $27cm/s$ it was demonstrated that the problems of hardware exhaustion, suboptimal results, and low efficiency as reported in the early studies of Hornby *et al.* [47], [48] (cf. section VII) could be overcome by narrowing down the search space, using a good walk engine, and well-selected initial walk parameters.

G. Lisbon 2004

Impressed by the practicality of UNSW's walk learning experiments in 2003 the teams from the University of Texas at Austin [64], [65], [108], the University of Newcastle in Australia [96], Carnegie Mellon University [21], Kyushu University [52], and the German Team [99] developed their own walk optimisation techniques. Some of the teams pushed the ERS-210s to walk with speeds of about $30cm/s$ and the new ERS-7s with speeds of about $40cm/s$. The preferred algorithm was simple reinforcement learning based on evolutionary algorithms. Most teams adopted a method for fitness evaluation as previously used by Hornby *et al.* [47], [48] or UNSW [20], [57], [100] which measured the time required by the robot to walk across a fixed distance between two coloured landmarks. The German Team [99] employed a different fitness evaluation which incorporated the robot paws' touch sensors to determine the time the robot's paws had ground contact while making a step as well as readings from the robot's acceleration sensors to minimise vibrations of the robot while walking. As documented by Kohl and Stone [64] most teams achieved comparable results and improved clearly over the maximum speed achieved by UNSW in the previous year [20], [57], [100].

Fidelman and Stone [31] built on the experience from walk learning and approached the task to let the ERS-7 learn to grasp the ball. The robot had to learn to approach the ball and to control it under its chin to have it ready for a kick. Signals from the infrared range sensor on the robot's chest were used to provide the learning algorithms with a binary reward signal (ball capture = 1; failure = 0). Fidelman and

Stone employed algorithms that previously had been used for walk learning and found that hill climbing, the amoeba algorithm, and policy gradient reinforcement learning lead to comparable results. They reported it took the robot about three hours or roughly 672 trials to increase an initial ball acquisition success rate of 36% to 64%. This study was a fully autonomous robot learning experiment where all training was done on the physical robot [31].

In contrast to some of the previously discussed studies which aimed at fully autonomous learning on the robot the UChile team developed a new AIBO simulator [124] which they used for learning of ball kicking behaviour [123] and walking [126]. In their 'Back to Reality' paradigm they combined and co-evolved their learning systems in the simulator and on the physical robot. By minimising the differences between fitness values obtained in the real and the virtual environment the simulator was adjusted to match the real environment as good as possible. In pilot experiments [126] improvements in walking speed could be achieved by using a combination of genetic search for the simulator parameters and policy gradient reinforcement learning for associated learning on the physical robot. The fastest learned walk had a speed close to $25cm/s$, which is clearly below the fastest walks achieved by some of the other teams in 2004. Zagal and Ruiz-del-Solar [124] claimed the reason for this suboptimal result may have been the restrictive design of their walk engine.

While ball, goal, and beacon recognition are essential skills for successful soccer play the task of opponent and team player recognition was neglected by many teams during the early years of the league. Members of the German team [122] implemented for test purposes a robot recognition method based on C4.5 decision tree learning [94]. Their approach consisted of image segmentation, attribute calculation, classification, and analysis. One complete processing cycle took only $27ms$ and could be run on the AIBO. The most time consuming step was attribute calculation which involved a slow iterative end-point algorithm [30]. The classification using decision trees was the fastest component [122].

VII. Machine Learning on AIBOs in General

Several universities and other research groups conducted studies involving machine learning on AIBOs which may not have been used in the competition code of the four-legged league but have value from a research point of view.

Among the first machine learning applications on the AIBO robots were studies by researchers from Sony Corporation who employed an evolutionary algorithm with crossover and mutation to develop gaits for a prototype of AIBO [47] and for the first consumer versions of AIBO [48]. Their walk engine had about 20 parameters which were evolved to obtain non-falling gaits which followed either a crawl, trot, or pace pattern. The experiments started from a hand developed crawl gait of about $6cm/s$ and reached $17cm/s$ for a pace and $10.8cm/s$ for a trot. The evolutionary algorithm had a population size of 30 and ran for 11 generations where each generation took about an hour to be processed on the physical robot [47]. Later experiments [48] extended the search space to find gaits that were effective on different types of surfaces. Hornby *et al* [47] reported that several parts of the robot had to be replaced repeatedly throughout the long walk learning experiments. This was possible because the researchers at Sony had sufficient access to new robot parts. However, the high hardware demands acted as a deterrent for researchers outside Sony and it would take about three years until the evolution of walk parameters for walking speed optimisation was successfully attempted by teams within the four legged league (cf., sections V-B, VI-F, and VI-G).

In 2002 Hardt and von Stryk [44] employed iterative numerical optimisation methods to minimise various performance and stability objectives for gait generation using a model of the AIBO robot based on kinematic and kinetic data provided by Sony. The model consisted of a 9-link tree-structure multibody system with central torso, a head at a fixed position, and four two-link legs. Hardt and von Stryk [44] mentioned that some of their simulations initially aimed at speeds of $67cm/s$ and possibly an implementation on the real robot at RoboCup 2002. However, as later developments on this topic would show robust walking with speeds of that magnitude were not realistic on the physical AIBO robot.

In the same year researchers at Essex [38] conducted a study on evolutionary gait generation on a physical AIBO robot. A genetic algorithm was employed to optimise 13 parameters of their locomotion module in 50 generations with a population size of 20. An overhead camera determined the extent of the robot's movement over five steps. Gyro sensor readings to determine walk stability were included with the movement measure into a fuzzy logic based fitness evaluation. The very low number of steps which were used for movement measurements protected the hardware but was not a precise speed measure. This, their large parameter space, and possible constraints of their walk engine may explain why they only achieved a maximum speed of $11.7cm/s$ [38] which is slow compared to other approaches at the same time (cf., sections V-B, VI-F, and VI-G).

In the initial years of the four-legged league another important task was to implement a vision system which achieved good results under constant and defined lighting conditions. After several teams had implemented vision systems with satisfactory performance it became apparent that most of them required long calibration times because they heavily relied on colour recognition and were extremely sensitive to small changes in lighting conditions. The task of colour constancy or illumination invariance, that is, to be able to compensate for changes in lighting conditions was recently approached by several groups associated with the four legged league.

Schulz and Fox [103] proposed a two-level Gaussian model where a switching Kalman filter estimated the lighting conditions at the upper level. The different filters were initialised using Gaussian priors. The latter were learned independently on collections of training images where images with similar lighting conditions were clustered using k-means clustering. Further a Rao-Blackwellised particle filter was employed to take the the robot's location in account by maintaining a joint posterior over robot positions and lighting conditions.

Sridharan and Stone [104] approached the colour constancy

TABLE III

OVERVIEW OF PUBLICATIONS WHICH REPORT ON THE USE OF DECISION TREES (DT), NEURAL NETWORKS AND SUPPORT VECTOR MACHINES (NN/SVM), EVOLUTIONARY COMPUTATION AND REINFORCEMENT LEARNING (EC/RL), OR OTHER MACHINE LEARNING METHODS ON AIBO ROBOTS. THE SECOND COLUMN (#TEAMS) IS THE NUMBER OF TEAMS IN THE FOUR-LEGGED LEAGUE AND THE LAST COLUMN (#PUB) INDICATES THE (APPROXIMATE) NUMBER OF PUBLICATIONS FOR EACH YEAR.

Year	#Teams	DT	NN/SVM	EC/RL	Other	#Pub
1998	3				[118] [119]	2
1999	9		[63]	**[47]**	[68]	3
2000	12	[77]	[50]	**[48]**	[45] [72]	5
2001	16	[6] [19]	[51] [49] **[69]**	[6] [39] [79]	**[69]** [88]	8
2002	19	[24] **[78]**		[26] **[37]** [40]	**[44]** **[56]** **[38]** [120]	9
2003	24	[5] [20]	[70] [96] [97]	[20] **[41]** [57] [96] [100] [125]	[7] **[55]** [106]	12
2004	24	**[114]** [122]	[43]	[21] [31] **[34]** [52] [64] [65] **[67]** [99] [108] [123] [126]	**[81]** **[85]** **[103]** **[104]** **[112]**	19

task by proposing a nearest neighbour algorithm where a set of hand labelled images is used to train colour cubes to map arrays of 128^3 pixel values to one of 10 relevant colours. With the KL-divergence as measure for comparing image distributions they proposed a method which allows the robot to recognise and adapt to three discrete illumination conditions: bright, intermediate, and dark.

Among the labs which installed an additional overhead camera for research purposes above the four-legged league soccer field was the group associated with the UChile four-legged team. The UChile group employed a delta-rule based on-line learning algorithm for behaviour learning with an AIBO robot [85]. The learned task was to move three randomly placed plastic cylinders to new positions on the soccer field until they formed the shape of a triangle.

In addition to the studies mentioned so far there were several other projects which demonstrated that machine learning on AIBOs was employed for a large variety of tasks outside the four-legged league competition. For example, Mitsunaga and Asada [78] employed decision trees for sensor space segmentation. Dynamic programming [113] was adopted for motion planning by Fukase *et al.* [37]. Learning of simple sensorimotor tasks using a convolutional neural network that automatically combines color, luminance, motion, and auditory information was presented by Lee and Seung [69]. Gu *et al.* [41] from Essex aimed at using a genetic algorithm and transfer from a simulation environment to learn ball chasing and position reaching behaviours. An unsupervised technique for the autonomous simultaneous calibration of action and sensor models was implemented on the ERS-7 by Stronger and Stone [112]. Vail and Veloso [114] used C4.5 for surface prediction and employed accelerometer data for for velocity prediction using a k-nearest neighbour algorithm [75]. Kwok and Fox [67] employed least square policy iteration for learning sensing strategies which led to improved goal scoring behaviour in the four-legged league domain. Some possibilities of using non-linear dimensionality reduction methods in the legged league domain were explored by Murch [81].

AIBOs became a popular platform for research in new aspects of robotics such as human–robot, animal–robot, or robot–robot interaction. Several of these projects have involved machine learning. Examples are robotic clicker training [56], experiments on neural learning for pointing gestures between two robots [43], or studies in curiosity-driven developmental robotics [55], [86], [87]. Researchers at Sony implemented AIBO's behaviour control architecture using probabilistic state machines whose probabilities are modified using reinforcement learning through interaction with the user [34].

VIII. DISCUSSION

The survey of Section VI revealed that the use of machine learning in the four-legged league was limited but growing. Table III presents a combined chronological overview of projects using machine learning on AIBO robots in general, and at the four-legged league competition. The references to general projects (from Section VII) are emphasised in boldface. Although a precise match of the year when a study was conducted and the year when the results were published was not always possible, a clear upwards trend of the use of machine learning on AIBOs is indicated in Table III. A dramatic increase in the use of machine learning can be observed in 2003. By then evolutionary [9], reinforcement learning algorithms [11], [113], or other simple search algorithms [91] were used for parameter optimisation and led to significant increases in the walking speed of the AIBO. Related approaches and spin-off studies by several competition teams and research groups then followed. These advancements in walking speed through machine learning were primarily achieved by the more successful and experienced teams of the four-legged league. Mutual inspiration of ideas and comparison of results between these research groups is documented in the associated publications, for example [57], [64], [65], [96], [99].

It is agreed by many researchers in the field that, given the challenges presented in Section II, the successful application of, for example, evolutionary or reinforcement learning methods on AIBO's to obtain results superior to handcoding, is a triumph. However, when compared with the very large amount of software and the vast array of algorithms produced by each team, machine learning still plays only a minor role. So far only a few machine learning techniques appear to be commonly used by the four-legged league's robot programmers. The most prominent among them are the C4.5 decision tree algorithm [94], neural networks, support vector machines, various simple clustering methods, and some basic evolutionary and reinforcement learning algorithms.

In the above-mentioned studies on walk learning, for example, it becomes apparent that most of the approaches employed very simple algorithms which are good for fast evaluation as this is important for experiments where large numbers of iterations can be performed. However, the task of walk learning with AIBOs only allows small numbers of iterations but there is plenty of time between the updates so that longer calculations with much more sophisticated algorithms could be performed. More sophisticated algorithms for this task, however, are not as well-known or as "practical" as the basic algorithms. The design and application of more advanced algorithms that address this issue might be a direction for future research.

This indicates that for the roboticists of the four-legged league the practicality of a machine learning method is an essential condition. "Practicality" means it has already been shown that the method works efficiently and a good implementation exists which is convenient to use and well-documented. A machine learning approach was typically only considered if it promised to significantly improve a white box approach, or in situations where the latter was not feasible.

The growing use of machine learning in the four-legged league demonstrates that RoboCup is progressing according to its mission to foster research through a robot soccer competition [60]. This is corroborated by the observation that there is an increasing number of projects using machine learning on AIBO robots which appear to be spin-off projects associated with the competitions at RoboCup (see Section VII). In contrast to the application oriented machine learning projects of the soccer competition the spin-off projects tended to take method oriented approaches and to investigate newer and more advanced machine learning methods or tasks.

IX. Conclusion

RoboCup is an exciting initiative which accelerates research into intelligent robotics, multi-agent systems, and robot learning. The four-legged league is important because it is the only real robot league at RoboCup where all teams employ the same prescribed hardware and therefore can focus on software development and algorithm design.

For an evaluation of the use of machine learning in the four-legged league and on AIBO robots it is appropriate to distinguish between application oriented projects of the competition teams, and method oriented studies of general research projects using AIBOs.

A chronological review showed that the number of publications per year on robot learning with AIBO robots increased from about 2 to 19 between 1998 and 2004. The use of machine learning in the four-legged league was initially limited but later it was growing rapidly and led to impressive results.

Recently Sony announced that they will cease the production of AIBO robots. Since AIBO robots have been such an excellent research platform, it is hoped that a suitable replacement can be found soon so that this fruitful initiative of software focused research on robot learning, as it was developing within the Sony four-legged league, can successfully be continued. Sony are to be commended for the significant contribution that AIBO robots have made to the development of machine intelligence research.

Acknowledgements

The first author would like to thank the RIKEN Brain Science Institute in Japan and AIFB at the University of Karlsruhe in Germany for hospitality during sabbatical. The authors are grateful to Rick Middleton and CDSC for their support. Finally the authors would like to thank the NUbots as well as the other four-legged teams at RoboCup whose projects were the basis of this article.

References

[1] "AIBO site at Sony corporation," http://www.sony.net/Products/aibo/, accessed 21. October 2005.

[2] "Federation International Robosoccer Association (FIRA)," http://www.fira.net/, accessed 1. June 2005.

[3] "RoboCup official site," http://www.robocup.org/, accessed 1. June 2005.

[4] "RoboCup special interest group (SIG) on multiagent learning," http://sserver.sourceforge.net/SIG-learn/, accessed 21. October 2005.

[5] H. L. Akın, M. K. Baloğlu, H. K. Bağcı, S. Bayhan, c. Meriçli, D. Poslu, O. Sever, O. T. Yıldız, S. Argirov, B. Marinov, P. Pavlova, N. Shakev, J. Tombakov, and A. Topalov, "Cerberus 2003 team report," Boğaziçi University and Technical University Sofia, Tech. Rep., 2003.

[6] H. Akın, A. Topalov, and O. Kaynak, "Cerberus 2001 Team Description," ser. LNCS, A. Birk, S. Coradeschi, and S. Tadokoro, Eds., vol. 2377. Springer, 2002, pp. 689–692.

[7] S. Anderson, K. Croaker, J. Dockter, V. Estivill-Castro, J. Fenwick, N. Lovell, and S. Seymon, "MiPAL Team Griffith - Summary of our technology and implementation," Griffith University, Tech. Rep., 2003.

[8] M. Asada and H. Kitano, Eds., *RoboCup-1998: Robot Soccer World Cup II*, ser. LNCS, vol. 1604. Springer, 1999.

[9] T. Baeck, D. Fogel, Z. Michalewicz, and S. Pidgeon, *Handbook of Evolutionary Computation*. IOP Publishing Ltd and Oxford University Press, 1997.

[10] A. G. Barto, R. S. Sutton, and C. W. Anderson, "Neuron-like adaptive elements that can solve difficult learning control problems," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 13, pp. 834–846, 1983.

[11] D. Bertsekas and J. Tsitsiklis, *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA, 1996.

[12] A. Birk, S. Coradeschi, and S. Tadokoro, Eds., *RoboCup-2001: Robot Soccer World Cup V*, ser. LNCS, vol. 2377. Springer, 2002.

[13] B. E. Boser, I. M. Guyon, and V. N. Vapnik, "A training algorithm for optimal margin classifiers," in *Proceedings of the 5th Annual ACM Workshop on Computational Learning Theory*, D. Haussler, Ed. Pittsburgh, PA: ACM Press, July 1992, pp. 144–152.

[14] M. Bowling, B. Browning, A. Chang, and M. Veloso, "Plays as team plans for coordination and adaptation," in *IJCAI Workshop on Issues in Designing Physical Agents for Dynamic Real-Time Environments: World modeling, planning, learning, and communicating. Acapulco, Mexico*, 2003.

[15] M. Bowling and M. Veloso, "Simultaneous adversarial multi-robot learning," in *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico*, 2003, pp. 699–704.

[16] J. Bruce, T. Balch, and M. Veloso, "Fast and inexpensive color image segmentation for interactive robots," in *Proceedings of the 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-2000)*, vol. 3, October 2000, pp. 2061–2066.

[17] G. Buskey, J. Roberts, and G. Wyeth, "A helicopter named Dolly—behavioural cloning for autonomous helicopter control," in *Proceedings of the Australasian Conference on Robotics and Automation*, W. Friedrich and P. Lim, Eds., Aukland, 2002.

[18] S. K. Chalup and C. L. Murch, "Machine learning in the four-legged league." in *3rd IFAC Symposium on Mechatronic Systems, September 6-8, 2004*, 2004.

[19] S. Chan, M. Sio, T. Volgelgesang, T. F. Yik, B. Hengst, S. B. Pham, and C. Sammut, "The UNSW RoboCup 2001 Sony Legged Robot League Team," ser. LNCS, A. Birk, S. Coradeschi, and S. Tadokoro, Eds., vol. 2377. Springer, 2002, pp. 39–45.

[20] J. Chen, E. Chung, R. Edwards, N. Wong, B. Hengst, C. Sammut, and W. Uther, "Rise of the AIBOS III - AIBO Revolutions," University of New South Wales, Tech. Rep., 2003.

[21] S. Chernova and M. Veloso, "An evolutionary approach to gait learning for four-legged robots," in *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-2004)*, 2004.

[22] J. Connell and S. Mahadevan, Eds., *Robot Learning*. Kluwer Academic Press, 1993.

[23] C. Cortes and V. Vapnik, "Support vector networks," *Machine Learning*, vol. 20, pp. 273–297, 1995.

[24] Z. Crisman, E. Curre, C. Kwok, L. Meyers, N. Ratliff, L. Tsybert, and D. Fox, "Team Description: UW Huskies-02," ser. LNCS, G. A. Kaminka, P. U. Lima, and R. Rojas, Eds., vol. 2752. Springer, 2002.

[25] N. Cristianini and J. Shawe-Taylor, *An Introduction to Support Vector Machines*. Cambridge University Press, 2000.

[26] I. Dahm and J. Ziegler, "Adaptive methods to improve self-localization in robot soccer," ser. LNCS, G. A. Kaminka, P. U. Lima, and R. Rojas, Eds., vol. 2752. Springer, 2002, pp. 393–408.

[27] J. Dalgliesh and M. Lawther, "Playing soccer with quadruped robots," Computer Engineering Thesis, 1999.

[28] J. Demiris and A. Birk, Eds., *Interdisciplinary Approaches to Robot Learning*. World Scientific, 2000.

[29] M. Dorigo, "Introduction to the special issue on learning autonomous robots," *IEEE Transactions on Sytems, Man, and Cybernetics*, vol. 26, pp. 361–364, June 1996.

[30] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*. John Wiley & Sons, 1973.

[31] P. Fidelman and P. Stone, "Learning ball acquisition on a physical robot," in *International Symposium on Robotics and Automation (ISRA), August 2004*, 2004.

[32] J. Franklin, T. Mitchell, and S. Thrun, "Robot learning," *Machine Learning Journal*, vol. 23, no. 2–3, 1996, special issue.

[33] M. Fujita, "AIBO: Toward the era of digital creatures," *The International Journal of Robotics Research*, vol. 20, pp. 781–794, 2001.

[34] ——, "On activating human communications with pet-type robot AIBO," *Proceedings of the IEEE*, vol. 92, no. 11, pp. 1804–1813, 2004.

[35] M. Fujita and K. Kageyama, "An open architecture for robot entertainment," in *Proceedings of the First International Conference on Autonomous Agents*, 1997, pp. 435–442.

[36] M. Fujita and H. Kitano, "Development of an autonomous quadruped robot for robot entertainment," *Autonomous Robots*, vol. 5, no. 1, pp. 7–18, 1998.

[37] T. Fukase, Y. Kobayashi, R. Ueda, T. Kawabe, and T. Arai, "Real-time decision making under uncertainty of self-localisation results," ser. LNCS, G. A. Kaminka, P. U. Lima, and R. Rojas, Eds., vol. 2752. Springer, 2002, pp. 375–383.

[38] D. Golubovic and H. Hu, "A hybrid evolutionary algorithm for gait generation of sony legged robots," in *Proceedings of the* 28^{th} *Annual Conference of the IEEE Industrial Electronics Society, Sevilla, Spain, November 5-8, 2002*, 2002.

[39] D. Gu and H. Hu, "Evolving fuzzy logic controllers for Sony legged robots," ser. LNCS, A. Birk, S. Coradeschi, and S. Tadokoro, Eds., vol. 2377. Springer, 2002, pp. 356–361.

[40] ——, "Reinforcement learning of fuzzy logic controller for quadruped walking robots," in *Proceedings 15th IFAC World Congress, Barcelona, Spain, July 21-26, 2002.*, 2002.

[41] D. Gu, H. Hu, J. Reynolds, and E. Tsang, "GA-based learning in behaviour based robotics," in *Proceedings of IEEE International Symposium on Computational Intelligence in Robotics and Automation, Kobe, Japan, 16-20 July 2003*, vol. 3, 2003, pp. 1521–1526.

[42] J.-S. Gutmann and D. Fox, "An experimental comparison of localization methods continued," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2002.

[43] V. V. Hafner and F. Kaplan, "Learning to interpret pointing gestures: Experiments with four-legged autonomous robots," in *NeuroBotics, Workshop Proceedings*, G. Palm and S. Wermter, Eds., 2004.

[44] M. Hardt and O. von Stryk, "The role of motion dynamics in the design, control and stability of bipedal and quadrupedal robots," ser. LNCS, G. A. Kaminka, P. U. Lima, and R. Rojas, Eds., vol. 2752. Springer, 2002, pp. 206–223.

[45] B. Hengst, D. Ibbotson, S. B. Pham, J. Dalgliesh, M. Lawther, P. Preston, and C. Sammut, "The UNSW RoboCup 2000 Sony Legged League Team," ser. LNCS, P. Stone, T. Balch, and G. Kraetzschmar, Eds., vol. 2019. Springer, 2001, pp. 64–75.

[46] B. Hengst, S. Pham, D. Ibbotson, and C. Sammut, "Omnidirectional locomotion for quadruped robots," ser. LNCS, A. Birk, S. Coradeschi, and S. Tadokoro, Eds., vol. 2377. Springer, 2002, pp. 368–373.

[47] G. S. Hornby, M. Fujita, S. Takamura, T. Yamamoto, and O. Hanagata, "Autonomous evolution of gaits with the Sony quadruped robot," in *Proceedings of 1999 Genetic and Evolutionary Computation Conference (GECCO), Vol. 2*, Banzhaf, Daida, Eiben, Garzon, Honavar, Jakiela, and Smith, Eds. Morgan Kaufmann, 1999, pp. 1297–1304.

[48] G. S. Hornby, S. Takamura, J. Yokono, O. Hanagata, T. Yamamoto, and M. Fujita, "Evolving robust gaits with AIBO," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA-2000)*, 2000, pp. 3040–3045.

[49] H. Hu, D. Gu, D. Golubovic, B. Li, and Z. Lio, "Essex Rovers 2001 Team Description," ser. LNCS, A. Birk, S. Coradeschi, and S. Tadokoro, Eds., vol. 2377. Springer, 2002, pp. 697–700.

[50] H. Hu and D. Gu, "A multi-agent system for cooperative quadruped walking robots," in *Proceedings IASTED International Conference Robotics and Applications, August 14–16, 2000 – Honolulu, Hawaii, USA*, 2000.

[51] ——, "Reactive behaviours and agent architecture for sony legged robots to play football," *International Journal of Industrial Robot*, vol. 28, no. 1, pp. 45–53, January 2001.

[52] J. Inoue, H. Aoyama, A. Ishino, and A. Shinohara, "Jolly pochie 2004 in the four legged robot league," Department of Informatics, Kyushu University 33, Fukuoka 812-8581, Japan, http://www.i.kyushu-u.ac.jp/JollyPochie/, Tech. Rep., 2004.

[53] A. Isaac and C. Sammut, "Goal-directed learning to fly," in *Proceedings of the Twentieth International Conference on Machine Learning*, T. Fawcett and N. Mishra, Eds., Washington, D.C., 2003, pp. 258–265.

[54] G. A. Kaminka, P. U. Lima, and R. Rojas, Eds., *RoboCup-2002: Robot Soccer World Cup VI*, ser. LNCS, vol. 2752. Springer, 2003.

[55] F. Kaplan and P.-Y. Oudeyer, "Motivational principles for visual know-how development," in *Proceedings of the 3rd Epigenetic Robotics workshop: Modeling cognitive development in robotic systems*, ser. Lund University Cognitive Studies, C. Prince, L. Berthouze, H. Kozima, D. Bullock, G. Stojanov, and C. Balkenius, Eds., vol. 101, 2003, pp. 72–80.

[56] F. Kaplan, P.-Y. Oudeyer, E. Kubinyi, and A. Miklosa, "Robotic clicker training," *Robotics and Autonomous Systems*, vol. 38, no. 3-4, pp. 197–206, 2002.

[57] M. S. Kim and W. Uther, "Automatic gait optimisation for quadruped robots," in *Australasian Conference on Robotics and Automation (ACRA'2003)*, 2003.

[58] H. Kitano, Ed., *RoboCup-1997: Robot Soccer World Cup I*, ser. LNCS, vol. 1395. Springer, 1998.

[59] H. Kitano, M. Fujita, S. Zrehen, and K. Kageyama, "Sony legged robot for RoboCup challenge," in *Proceedings of the 1998 IEEE International Conference on Robotics and Automation.*, vol. 3, 16-20 May 1998, pp. 2605–2612.

[60] H. Kitano, M. Asada, Y. Kuiyoshi, I. Noda, E. Osawa, and H. Matsubara, "Robocup: A challenge problem for AI and robotics," ser. LNCS, H. Kitano, Ed., vol. 1395. Springer, 1997.

[61] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, E. Osawa, and H. Matsubara, "Robocup: A challenge problem for AI," *AI Magazine*, vol. 18, no. 1, pp. 73–85, Spring 1997.

[62] A. Kleiner, M. Dietl, and B. Nebel, "Towards a life-long learning soccer agent," ser. LNCS, G. A. Kaminka, P. U. Lima, and R. Rojas, Eds., vol. 2752. Springer, 2002.

[63] Y. Kobayashi and H. Yuasa, "Team ARAIBO," ser. LNCS, M. Veloso, E. Pagello, and H. Kitano, Eds., vol. 1856. Springer, 2000, pp. 758–761.

[64] N. Kohl and P. Stone, "Machine learning for fast quadrupedal locomotion," in *The Nineteenth National Conference on Artificial Intelligence*, 2004.

[65] ——, “Policy gradient reinforcement learning for fast quadrupedal locomotion,” in *IEEE International Conference on Robotics and Automation*, 2004.

[66] C. Kwok and D. Fox, “Map-based multi model tracking of a moving object,” ser. LNCS, D. Nardi, M. Riedmiller, C. Sammut, and J. Santos-Victor, Eds., vol. 3276. Springer, 2004.

[67] ——, “Reinforcement learning for sensing strategies,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-04)*, 2004.

[68] M. Lawther and J. Dalgliesh, “UNSW United,” ser. LNCS, M. Veloso, E. Pagello, and H. Kitano, Eds., vol. 1856. Springer, 2000, pp. 788–791.

[69] D. D. Lee and H. S. Seung, “Biologically inspired computation and learning in sensorimotor systems,” in *Applications and Science of Neural Networks, Fuzzy Systems, and Evolutionary Computation IV*, B. Bosacchi, D. B. Fogel, and J. C. Bezdek, Eds., vol. 4479, no. 1. SPIE, 2001, pp. 4–11.

[70] B. Li, H. Hu, and L. Spacek, “An adaptive color segmentation algorithm for sony legged robots,” in *Proceedings of the 21st IASTED International Conference on Applied Informatics, February 1—13, 2003, Insbruck, Austria*, 2003, pp. 126–131.

[71] S. Mahadevan, “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)*, 1996.

[72] G. Marceau, “The McGill’s Red Dogs Legged League System,” ser. LNCS, P. Stone, T. Balch, and G. Kraetzschmar, Eds., vol. 2019. Springer, 2001, pp. 627–630.

[73] M. Mataric and D. Cliff, “Challenges in evolving controllers for physical robots,” *Robotics and Autonomous Systems*, vol. 19, no. 1, pp. 67–83, 1996.

[74] R. H. Middleton, M. Freeston, and L. McNeill, “An application of the extended kalman filter to robot soccer localisation and world modelling,” in *IFAC Symposium on Mechatronic Systems*, 2004.

[75] M. Mitchell, *An Introduction to Genetic Algorithms*. MIT Press, 1996.

[76] T. Mitchell, *Machine Learning*. McGraw Hill, 1997.

[77] M. Mitsunaga and M. Asada, “Observation strategy for decision making based on information criterion.” ser. LNCS, P. Stone, T. Balch, and G. Kraetzschmar, Eds., vol. 2019. Springer, 2001, pp. 189–198.

[78] ——, “Visual attention control by sensor space segmentation for a small quadruped robot based on information criterion,” ser. LNCS, A. Birk, S. Coradeschi, and S. Tadokoro, Eds., vol. 2377. Springer, 2002, pp. 154–163.

[79] N. Mitsunaga, Y. Nagai, T. Ishida, T. Izumi, and M. Asada, “BabyTigers 2001: Osaka Legged Robot Team,” ser. LNCS, A. Birk, S. Coradeschi, and S. Tadokoro, Eds., vol. 2377. Springer, 2002, pp. 685–688.

[80] N. Mitsunaga and M. Asada, “BabyTigers 1999: Osaka Legged Robot Team,” ser. LNCS, M. Veloso, E. Pagello, and H. Kitano, Eds., vol. 1856. Springer, 2000, pp. 762–765.

[81] C. L. Murch, “Dimensionality reduction on AIBO robots,” Honours Thesis, School of Electrical Engineering and Computer Science, The University of Newcastle, Australia, November 2004.

[82] D. Nardi, M. Riedmiller, C. Sammut, and J. Santos-Victor, Eds., *RoboCup-2004: Robot Soccer World Cup VIII*, ser. LNCS, vol. 3276. Springer, 2005.

[83] A. Y. Ng, H. J. Kim, M. I. Jordan, and S. Sastry, “Autonomous helicopter flight via reinforcement learning,” in *Advances in Neural Information Processing Systems 16*, S. Thrun, L. Saul, and B. Schölkopf, Eds. Cambridge, MA: MIT Press, 2003.

[84] NIPS conference, “Advances in neural information processing systems,” on-line at http://nips.cc or printed version by MIT Press.

[85] X. Olivares, J. Ruiz-del-Solar, and J. C. Zagal, “On-line learning of an object manipulation behaviour for legged robots,” in *Proceedings of the 1st IEEE Latin American Robotics Symposium—LARS 2004, Mexico City, Mexico, October 28 - 29, 2004*, 2004, pp. 42–47.

[86] P.-Y. Oudeyer and F. Kaplan, “Intelligent adaptive curiosity: a source of self-development,” in *Proceedings of the 4th Epigenetic robotic workshop, Genoa*, 2004.

[87] P.-Y. Oudeyer, F. Kaplan, V. Hafner, and W. A., “The playground experiment: Task-independent development of a curious robot,” in *proceedings of the AAAI Spring Symposium Workshop on Developmental Robotics*, 2005.

[88] S. Pham, B. Hengst, D. Ibbotson, and C. Sammut, “Stochastic gradient descent localisation in quadruped robots,” ser. LNCS, A. Birk, S. Coradeschi, and S. Tadokoro, Eds., vol. 2377. Springer, 2002, pp. 447–452.

[89] T. Poggio and S. Smale, “The mathematics of learning: Dealing with data,” *Notices of the American Mathematical Society*, vol. 50, no. 5, pp. 537–544, 2003.

[90] D. Polani, B. Browning, A. Bonarini, and K. Yoshida, Eds., *RoboCup 2003: Robot Soccer World Cup VII*, ser. LNCS, vol. 3020. Springer, 2004.

[91] W. Press, S. Teukolsky, V. W.T., and B. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed. Cambridge University Press, 1995.

[92] F. K. H. Quek, “An algorithm for the rapid computation of boundaries of run-length encoded regions,” *Pattern Recognition Journal*, vol. 33, no. 10, pp. 1637–1649, October 2000.

[93] J. R. Quinlan, “Discovering rules from large collections of examples: A case study,” in *Expert Systems in the Microelectronic Age*, D. Michie, Ed. Univerisity Press, Edinburgh, Scotland, 1979.

[94] ——, *C4.5: Programs for Machine Learning*. Morgan Kaufmann, 1993.

[95] M. J. Quinlan, C. Murch, R. H. Middleton, and S. K. Chalup, “Traction monitoring for collision detection with legged robots,” ser. LNAI, D. Polani, B. Browning, A. Bonarini, and K. Yoshida, Eds., vol. 3020. Springer-Verlag, 2004.

[96] M. J. Quinlan, S. K. Chalup, and R. H. Middleton, “Techniques for improving vision and locomotion on the Sony AIBO robot,” in *Australasian Conference on Robotics and Automation (ACRA’2003)*, 2003.

[97] ——, “Application of SVMs for colour classification and collision detection with AIBO robots,” in *Advances in Neural Information Processing Systems 16 (NIPS’2003)*, S. Thrun, L. Saul, and B. Schölkopf, Eds. Cambridge, MA: MIT Press, 2004.

[98] M. Riedmiller and A. Merke, “Using machine learning techniques in complex multi-agent domains,” in *Perspectives on Adaptivity and Learning*, ser. LNCS, I. Stamatescu, W. Menzel, M. Richter, and U. Ratsch, Eds. Springer, 2002.

[99] T. Röfer, “Evolutionary gait-optimization using a fitness function based on proprioception,” ser. LNCS, D. Nardi, M. Riedmiller, C. Sammut, and J. Santos-Victor, Eds., vol. 3276. Springer, 2004.

[100] C. Sammut, “Robot soccer: Science or just fun and games ?” in *AI 2003: Advances in Artificial Intelligence 16th Australian Conference on AI*, ser. LNAI, T. D. Gedeon and L. C. C. Fung, Eds., vol. 2903, 2003, pp. 12–23.

[101] B. Schölkopf and A. J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, 2002.

[102] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson, “Estimating the support of a high-dimensional distribution,” *Neural Computation*, vol. 13, pp. 1443–1471, 2001.

[103] D. Schulz and D. Fox, “Bayesian color estimation for adaptive vision-based robot localization,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-04)*, 2004.

[104] M. Sridharan and P. Stone, “Towards illumination invariance in the legged league,” in *RoboCup-2004: Robot Soccer World Cup VIII*, ser. LNCS, D. Nardi, M. Riedmiller, and C. Sammut, Eds., vol. 3276. Berlin: Springer Verlag, 2005, pp. 196–208.

[105] P. Stone, T. Balch, and G. Kraetzschmar, Eds., *RoboCup-2000: Robot Soccer World Cup IV*, ser. LNCS, vol. 2019. Springer, 2001.

[106] P. Stone, K. Dresner, S. Erdoğan, P. Fidelman, N. Jong, N. Kohl, G. Kuhlmann, E. Lin, M. Sridharan, D. Stronger, and G. Hariharan, “UT Austin Villa 2003: A New RoboCup Four-Legged Team,” University of Texas at Austin, Tech. Rep., 2003.

[107] P. Stone, “Multiagent competitions and research: Lessons from RoboCup and TAC,” ser. LNCS, G. A. Kaminka, P. U. Lima, and R. Rojas, Eds., vol. 2752. Springer, 2002, pp. 224–237.

[108] P. Stone, K. Dresner, P. Fidelman, N. K. Jong, N. Kohl, G. Kuhlmann, M. Sridharan, and D. Stronger, “The UT Austin Villa 2004 RoboCup four-legged team: Coming of age,” The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, Tech. Rep. UT-AI-TR-04-313, October 2004.

[109] P. Stone, R. S. Sutton, and G. Kuhlmann, “Reinforcement learning for RoboCup-soccer keepaway,” *Adaptive Behavior*, vol. 13, no. 3, pp. 165–188, 2005.

[110] P. Stone and M. Veloso, “Multiagent systems: A survey from a machine learning perspective,” *Autonomous Robots*, vol. 8, no. 3, pp. 345–383, 2000.

[111] ——, “A survey of multiagent systems and multirobot systems,” in *Robot Teams: From Diversity to Polymorphism*, T. Balch and L. E. Parker, Eds., 2002, ch. 3, pp. 37–92.

[112] D. Stronger and P. Stone, "Simultaneous calibration of action and sensor models on a mobile robot," in *2004 International Symposium on Robotics and Automation (ISRA)*, August 2004.
[113] R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*. MIT Press, 1998.
[114] D. Vail and M. Veloso, "Learning from accelerometer data on a legged robot," in *Proceedings of the 5th IFAC/EURON Symposium on Intelligent Autonomous Vehicles (IAV2004), Lisbon, Portugal, July 2004*, 2004.
[115] V. N. Vapnik, *The Nature of Statistical Learning Theory*. Springer–Verlag, 1995.
[116] ——, *Statistical Learning Theory*. New York: Wiley, 1998.
[117] M. Veloso, E. Pagello, and H. Kitano, Eds., *RoboCup-1999: Robot Soccer World Cup III*, ser. LNCS, vol. 1856. Springer, 2000.
[118] M. Veloso, W. Uther, M. Fujita, M. Asada, and K. H., "Playing soccer with legged robots," in *Proceedings of the International Conference on Intelligent Robots and Systems Conference (IROS-98)*, 1998, pp. 437–442.
[119] M. Veloso and W. Uther, "The CMTrio-98 Sony legged robot team," ser. LNCS, M. Asada and H. Kitano, Eds., vol. 1604. Springer, 1999, pp. 491–497.
[120] Z. Wang, J. Wong, T. Tam, B. Leung, M. S. Kim, J. Brooks, A. Chang, and N. Von Huben, "rUNSWift UNSW RoboCup2002 Sony Legged League Team," University of New South Wales, Tech. Rep., 2002.
[121] T. Weigel, J.-S. Gutman, M. Dietl, A. Kleiner, and B. Nebel, "CS Freiburg: Coordinating robots for successful soccer playing," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 685–699, 2002.
[122] D. Wilking and T. Röfer, "Realtime object recognition using decision tree learning," ser. LNCS, D. Nardi, M. Riedmiller, C. Sammut, and J. Santos-Victor, Eds., vol. 3276. Springer, 2004.
[123] J. C. Zagal and J. Ruiz-del-Solar, "Learning to kick the ball using back to reality," ser. LNCS, D. Nardi, M. Riedmiller, C. Sammut, and J. Santos-Victor, Eds., vol. 3276. Springer, 2005.
[124] ——, "UCHILSIM: A dynamically and visually realistic simulator for the robocup four legged league," ser. LNCS, D. Nardi, M. Riedmiller, C. Sammut, and J. Santos-Victor, Eds., vol. 3276. Springer, 2005, pp. 34–45.
[125] J. C. Zagal, J. Ruiz-del-Solar, P. Guerrero, and R. Palma, "Evolving visual object recognition for legged robots," ser. LNCS, D. Polani, B. Browning, A. Bonarini, and K. Yoshida, Eds., vol. 3020. Springer, 2004.
[126] J. C. Zagal, J. Ruiz-del-Solar, and P. Vallejos, "Back to reality: Crossing the reality gap in evolutionary robotics," in *Proceedings of the IAV 2004, 5th IFAC Symposium on Intelligent Autonomous Vehicles*, 2004.

Stephan K. Chalup is senior lecturer in computer science and software engineering at the University of Newcastle in Australia. His Ph.D. was awarded from Queensland University of Technology. He received his diploma in mathematics with biology from the University of Heidelberg in Germany in 1996. He is leader of the Interdisciplinary Machine Learning Research Group (IMLRG) and one of the academic chief investigators in the Newcastle Robotics Laboratory. His research interests include autonomous robots, computational neuroscience, kernel machines, and manifold learning.

Craig L. Murch was born in Chicago, USA. He received the degrees of B.CompSc and B.Math in 2002 and 2003 respectively from the University of Newcastle in Australia, and has just completed honours in computer science at the same institution. His research interests include dimensionality reduction and manifold learning.

Michael J. Quinlan was born in Newcastle, Australia. He received a B.E. degree in software engineering in 2000 from the University of Newcastle, Newcastle, Australia, where he is currently working towards a Ph.D. degree. His research interests include quadruped locomotion, support vector machines and robotic applications of machine learning.