The NUbots' Team Description for 2003

Stephan K. Chalup, Oliver J. Coleman, Michaela N. Freeston,
Richard H. Middleton, Craig L. Murch, Michael J. Quinlan,
Christopher J. Seysener, and Graham D. Shanks

School of Electrical Engineering & Computer Science
The University of Newcastle, Callaghan 2308, Australia
{chalup,rick}@eecs.newcastle.edu.au
WWW home page: http://robots.newcastle.edu.au

Abstract. The NUbots are the University of Newcastle's Legged-league RoboCup team. They achieved third place at RoboCup in 2002 and 2003. The team consists of students and academics from two disciplines, Electrical Engineering and Computer Science & Software Engineering. The present paper provides information about the team structure, the research interests of the team members, and the research and study environment which hosts the NUbots. A concise overview about the NUbots' approach to the RoboCup 2003 competition is included as well.

1 Introduction

The University of Newcastle, Australia has competed in RoboCup since 2001. The first team consisted of two undergraduate students who won the world title of RoboCup Junior 2001 in Seattle. In 2002 the NUbot team was formed and entered the Sony Legged League at RoboCup 2002 in Fukuoka where the new team instantly achieved a third place. In 2003 the NUbots were able to repeat their success and again came third at RoboCup in Padova. The 2003 NUbot team had improved its techniques and strength over 2002 and was only beaten in the semi finals by the team of the University of Pennsylvania 3-4 on penalties. The playoff for third place saw the NUbots beat last year's world champions Carnegie Mellon 4-1. In the 2003 competition the NUbots achieved the most goals for (83), and fewest goals against (3). It was noticed that they matched the previous record for the highest score in a legged league game at RoboCup by winning a match 16-0 against one of the other competitors.

2 Background of the NUbot Team Members

The 2003 NUbot team consists of five student developers and two academics. The main part of coding and strategy development is done by the student team:

– Michaela Freeston (Computer Engineering undergraduate)
– Craig Murch (Mathematics / Computer Science undergraduate)
– Michael Quinlan (PhD student)

- Chris Seysener (Computer Engineering / Computer Science undergraduate)
- Graham Shanks (Computer Science undergraduate)

Oliver Coleman, a Computer Science Honours student, worked separately on a simulated solution for the obstacle avoidance challenge. Background work on vision was done by Jared Bunting and Will McMahan during the second half of 2003. All developers were students in the School of Electrical Engineering & Computer Science at the University of Newcastle. Management, guidance and research supervision is responsibility of the two academic NUbot team members:

Prof. Rick Middleton (Team Leader) has published research results in a range of areas including electric machine control, adaptive control, robot control, digital control systems theory using delta operators, multirate and sampled-data systems, performance limitations in feedback control systems (including multivariable and nonlinear systems), metal rolling control problems, robotics. He is co-author of the text “Digital Control and Estimation: A Unified Approach” (Prentice-Hall). He has been involved in industrial applications of systems and control to radio astronomy, satellite tracking, metals processing industries, power electronic inverter controls and various applications of Kalman filtering. He has served as an associate editor of both the IEEE Transactions on Automatic Control and the IEEE Transaction on Control System Technology. He is an Associate Editor of Automatica and is Director of the Centre for Integrated Dynamics and Control (A Commonwealth Special Research Centre).

Stephan Chalup, Ph.D. is lecturer in Computer Science and Software Engineering. His background is machine learning, mathematics and neurobiology. He has published in the areas of neural networks, evolutionary computation, incremental learning, learning of formal languages, and robotics. He is supervisor of several research students and leader of the Interdisciplinary Machine Learning Research Group (IMLRG) at the University of Newcastle, Australia.

3 Related Research Concentrations

A number of researchers and research concentrations at the University of Newcastle are affiliated with the NUbot team. Currently these are primarily the CIDAC centre and the IMLRG research group. However, there is growing support from researchers in embedded systems, statistics, and theoretical computer science.

The *Centre for Integrated Dynamics and Control (CIDAC)* is a Special Research Centre funded by the Australian Government and linked to the School of Electrical Engineering and Computer Science at the University of Newcastle, Australia. The Centre provides significant industrial and manufacturing performance advances by working on approaches to control and scheduling. These approaches aim to unify the use of disparate technologies, namely, mathematical modelling through to computer systems, electromechanical machinery, scheduling systems and chemical processing. This will bring about an increase in the performance of industry in key areas including product quality, plant efficiency, safety, productivity, waste minimisation, pollution control and operational flexibility. For more details see `http://www.ee.newcastle.edu.au/cidac/`

The *Interdisciplinary Machine Learning Research Group (IMLRG)* is one of the research groups in the Discipline of Computer Science and Software Engineering at the University of Newcastle. It investigates different aspects of machine learning and data mining in theory, experiments and applications. Particular emphasis is put on interdisciplinary projects. For active members IMLRG is a forum for exchange and update of information about new literature, research tools and techniques. Some of IMLRG's research areas are: Data mining, machine learning, robotics, RoboCup, control and learning, neuroscience, bioinformatics, evolutionary computation, reinforcement learning, support vector machines. For more details see `http://www.cs.newcastle.edu.au/Research/IMLRG/index.html`

4 Research Interests

In 2003 the NUbots were working on several research projects related to robotics, control, and machine learning. The following subsections provide brief summaries of some of the NUbots' research projects and links to corresponding publications.

4.1 Vision

Vision is one of the major research areas of the NUbots. Several subtopics have been investigated including object recognition, horizon determination, edge detection, and colour classification using ellipse fitting, convex optimization and kernel machines. Publications are available e.g. from [8, 2, 1, 11, 10, 12]. In the future this research area might extend towards efficient schemes for edge detection and also for Bayesian approaches to video processing of image streams. We plan to look into applications of image processing to natural lighting conditions and investigate computationally efficient geometric circle fitting algorithms.

4.2 Localisation, Kalman Filters

Research on the topic of localisation focused on Bayesian approaches to robot localisation including Extended Kalman Filters and particle filter based methods. We are particularly interested in further modifications of the Kalman Filter to handle non-ideal information from vision. For information about our current approach see [6, 5, 2, 1].

4.3 Locomotion, Gait Optimisation

We analysed existing gaits of the robots and included low level controller parameters in the NUbots' locomotion engine. The task of gait optimisation involves learning in a poorly structured high dimensional parameter space. For this purpose we developed and tested different optimization schemes based on evolutionary computation [10].

4.4 Traction Monitoring

Methods to monitor traction measures are developed and employed for collision detection, to increase the speed of the robots, and to find a good strategy to deal with situations where the legs of two robots get entangled (*leg-lock*) [9]. The techniques used are examples of applications of fault detection ideas, which may further find use in monitoring other collisions and unusual situations.

4.5 Navigation and Obstacle Avoidance

The suitability of different path planning and obstacle avoidance techniques is compared in simulated environments to real-world environments on the task of the RoboCup Legged League 2003 Obstacle Avoidance Challenge. We investigate how the perceptions of the agents in the simulated environment and real-world environment impact on the choice of obstacle avoidance and path planning methods used for the task [4].

5 Overview of the Approach Taken in the 2003 Legged-League Competition

The NUbots' software system at RoboCup 2003 was a substantially revised and extended version of the 2002 system [2]. Technical details about software and concepts employed by the 2003 NUbot team are available in the NUbots' 2003 team report [1].

While the architecture of the 2002 software system consisted of many objects in 2003 all OPEN-R objects were merged into a single binary. An overview of the underlying conceptual system modules is given in Fig. 1. The main modules address vision, localisation, locomotion, and behaviour.

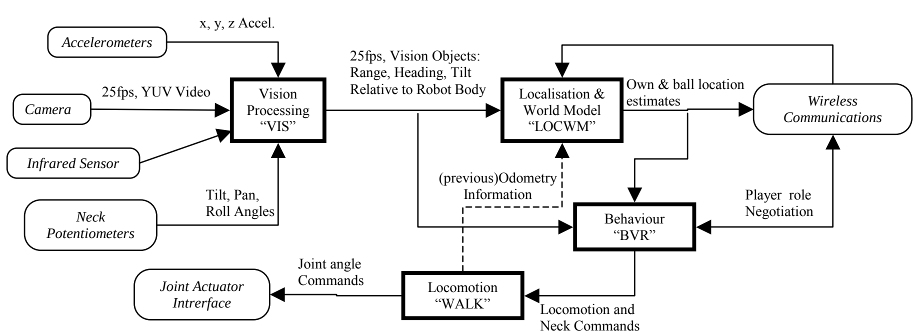

Fig. 1. Overview about the main modules and their connections in the NUbots' software system.

The main improvements over the 2002 approach were:

1. Better vision processing (with debugging, additional sanity checking and a number of novel features).
2. An integrated and enhanced Extended Kalman Filter for localisation and world modelling.
3. Increased walking speed and agility of the robots.
4. Improved behaviours including kick selections, position negotiation using wireless, and ball 'saves'.
5. For the 2003 system two helpful tools were employed:
   - The use of configuration files allowed parameter adjustments without recompiling or rebooting.
   - Wireless network debugging used a newly built-in telnet function.

5.1 Vision

The visual processing subsystem, is the result of a cumulative effort of a number of members of the NUbot team over the past two years [12, 1, 2, 8]. This system, like many others on the robot, is simple to explain yet complicated in its implementation. Humans process images with apparent ease, quickly filtering out useless information, identifying objects based on their shape and colour and then passing this information on to the brain for processing. When implemented on a robot this process is not as easy as it initially appears. The design of the vision system included an analysis of the camera's characteristics, including a response analysis and an image deviation/noise level analysis (for details see [1]). Image processing was performed in four steps:

- Colour Classification
- Run Length Encoding
- Blob Formation
- Object Recognition

Some of the major new features in the vision system for 2003 focused on improving the distance and heading data of existing data provided to the localisation module, while the second considered the provision of additional information to the localisation module, such as information about the corners of the penalty box [12].

Field and Sideline Detection Whilst determining the distance and heading to a beacon or goal is perhaps the easiest and most convenient way of performing localisation, problems with not observing beacons for significant periods of time have prompted research into finding additional points on the field from which to localise. One such idea was the use of the field and sideline information within the image to provide additional information about the robot's placement on the field. Under the implementation discussed in the report, the key to determining the presence of field and sidelines lies in the system's ability to discern field-green/boundary-white transitions in the classified image. Due to the possibility of multiple disjoint lines being present in the image, clustering algorithms were developed to group points identified in to field and sidelines. Horizon lines generated from the tilt, pan and roll of the robots body and head were used to filter parts of the image where transitions could not occur.

Penalty Box Corner Detection Goalkeepers often end up in the back corner of the goal unable to move. In addressing the problem of poor goalkeeper localisation, it was decided that it would be highly beneficial if additional information from which to localise was available. The solution was to use field line data and more specifically the placement of the penalty box corners within the image. The research completed focused on the fitting lines of two lines to the field line points already discovered and used a simple Hill Climbing algorithm to search in the direction that minimised the residual of two lines fitted to the dataset. The point that minimises the residual of the system is taken as the corner point. This algorithm results in an extremely efficient solution to the problem of finding a corner in a system of ordered points. The distance to the point identified is calculated based on the tilt of the robot and the height of the robot above the ground.

Centre Circle Detection The identification of the centre circle within the image provides yet another point from which the robot can localise. As the centre circle closely resembles a field line, it made sense that the detection of the centre circle should some how be related to the detection of field lines within the image. The presence of the centre circle is thus determined by the existence of a special case of field line detection where three or more sets of field line transitions are present in the same column. While discovery of the Centre Circle is only possible from certain positions upon the field, it does provide a foundation from which further work can be performed. Future work may include research into the fitting of circles or ellipses to determine accurately the presence and parameters of the circle.

Centre Corner Point Detection Another area in which field and side line information was used was in the identification of the two points where the centre field line intersects with the sideline. The detection of these points requires that a field line must necessarily intersect with a sideline and must not occur in the same image as a goal. On checking that the field line is sufficiently close to the sideline, the intersection of the field and sideline is taken as the Centre Corner point. The use of additional information such as the Centre Corner points greatly enhances the robot's ability to localise both with and without beacons on the field. The success of the Edge Recognition system implemented on the robot when combined with the 2003 Localisation subsystem was proven in the localisation without beacons challenge at RoboCup 2003 where a robot was able to move to five fixed positions on the field using only the goals and data provided by the edge detection system.

Ellipse and Circle Fitting to the Ball The process of fitting rectangular blobs to objects within the image, whilst convenient to code and adequate in practice, complicates the process of determining specific details about an object. Due to the significant loss of precision in fitting a rectangle to the circular shape of the ball, it was decided that some form of alternate algorithm should be considered. In ball recognition, it was decided that points located on the perimeter of the ball should be identified and a circle fitted through these points in order to make the system more tolerant of obstructions such as another robot or the edge of the image. Having reviewed the results of Chernov and Lesort [3] it was decided that a combination of both geometric and algebraic methods should be used to provide a robust solution to the fitting of a circle through the points identified. The implementation of the Levenberg-Marquardt fitting algorithm provided significantly better results than the 2002 algorithm, which calculated distance based solely on the height of the blob. This distance has proven sufficiently accurate for use in the behaviour model in deciding which robot should chase the ball.

Infrared Distance Detection Whilst the distances determined from the image are generally accurate, their accuracy can suffer when the height of the object is affected by an obstruction such as another robot. The infrared range sensor returns distances to objects where the measured value is less than 900mm. Tests were performed to identify the minimum area that an object blob needed to cover before its distance could be accurately measured by the infrared sensor. In some cases object specific coding was required to compensate for the irregular shape of objects such as the circular shape of the ball.

Ball on Wall Detection The presence of the ball adjacent to a wall can cause a number of problems such as potentially pushing the ball the full length of the field towards the robot's own goal. The process of detecting the ball next to the wall is reasonably simple. Whenever a ball is located on the wall, the colour on one side of the ball is white and the other field green. The process of checking for the wall is therefore to search left and right of the ball for a white/green or green/white combination. Having detected the presence of the ball on the wall, the robots behaviour is altered to force the robot to walk around the ball before kicking.

5.2 Localisation

The NUbots system of localisation and world modelling [5, 6] is based on an Extended Kalman Filter (EKF). An integrated, 5th order model for the robots own location and the ball location is used.

We included in this year's code sensing of the perceived angle between two objects; rather than treating each object in a vision frame independently. We also improved the mechanism for constraining the robot's localisation estimate to the field of play to include the area around the goal mouth. Previously, this area was implicitly excluded as being beyond the ‘goal line’.

The system includes an algorithm for dealing with both the lost robot problem, (or the ‘kidnapped robot problem’) and measurement outliers (or spurious measurements). We note that the lost robot problem is a modern form of a long studied problem in detecting and tracking jumps in an estimation setting.

The method recognises several successive (or almost successive) bad (i.e conflicting with the current model) vision measurements to distinct object types. By choosing appropriate threshold values for various constants (such as how many distinct object types had to be seen, how bad the measurements had to be, etc.), the robot is able to detect that it has been manually picked up and moved almost without fail. Also, where bad measurements were on a single object type, this was rejected as an outlier and not used in the model (for example if a beacon is removed and deliberately misplaced on the field). As well as detecting manual robot repositioning, the detection system allows the robot to quickly regain its bearings after receiving large amounts of erroneous odometry data (as in a severe collision with another robot).

An obvious technique to employ once a lost robot determination has been made is to reset the Kalman Filter uncertainty values. This allows the robot to very rapidly move to its new position without compromising localisation performance in normal situations.

The localisation challenge (in which all landmarks except the goals were removed) was performed with essentially no changes to the regular game code for localisation, and we appeared to be close to all required target points (though not close enough to score full points).

5.3 Locomotion

The NUbots’ new locomotion engine is inspired by the concept of parameterised walk (Pwalk) which was developed by the team of the University of New South Wales (UNSW) [7]. In depth empirical testing of different parameter configurations and relations which were obtained via inverse kinematics improved the walking, turning, and chasing speed of the robots and the effectiveness of their kicks.

The all new 2003 locomotion engine initially made use of an elliptical locus. It allowed considerably higher speed than the 2002 module. Unfortunately, despite months of tweaking, the locomotion module underachieved in the 2003 Australian Open in May. Much concerted experimentation (and the knowledge that a failure to improve would lead to a certain loss) resulted in a trapezoid based walk which was considerably faster again than the elliptical one. Additionally, reworked turning mathematics vastly improved the chasing ability of the robots.

The NUbot kicking system allows for quicker kick development and more flexible kick execution. Kicks are loaded from a file (rather than compiled into code like many teams), and this file can be reloaded at any time. This allows new kicks to be developed extremely rapidly. Interpolation was added in order to better tweak the speed of different parts of a kick motion. Interpolation was also added between certain motion transitions in order to make the robot move more naturally.

5.4 Behaviour

High level behaviour is responsible for the overall team strategy adopted by the NUbots. It is concerned with matters such as deciding which point on the field the robot should move to, and if/when the robot should be chasing the ball. With the introduction of wireless communication, high level behaviour has also become responsible for coordinating the actions of multiple robots in an efficient manner.

The main problems with last year's behaviour system were a direct result of its rigidity. Robot positions were static. The robot assigned to chase the ball was based solely on static regions regardless of whether they could see the ball, and indeed, regardless of their distance to the ball. As a result, we have tried to make our 2003 behaviour as flexible as possible.

One instance of applying this philosophy in the NUbot behaviour is that the closest robot to the ball will attack it. Each robot sends its perceived distance to the ball over the wireless network. Other robots will move to an area on the field determined by the captain. It is possible to change which player acts as captain at run-time, but the goalkeeper is typically assigned the captain role.

The captain assigns the chasing robot to the area containing the ball, while the other robots are assigned to other field areas based on their distances to them. These decisions are made using the captains global world model, although the determination of which robot is currently chasing is made based on wirelessly communicated ball distances. The set of available areas for robots to move to is defined based on the position allocated to robot chasing the ball. If the robot chasing the ball is in the opposition half, there will be two attackers and one defender. If the chasing robot is in the defensive half, there will be two defenders and one attacker.

It is important to realise that there is no explicit negotiation taking place. Since all the robots are equally capable of determining which robot is the optimal chaser, there is no need for an explicit decision making process to take place over the wireless network. This helps considerably in dealing with the latency induced by wireless communication on the AIBO. The result is very effective dynamic play in the midfield, while regions are retained only for the penalty box (as required by RoboCup rules). Robots not currently chasing will look for the ball in the direction indicated by their shared world model. They will also move around within their assigned position area based on the position of the ball in order to maximise their chance of gaining possession.

6 Study Options in Robotics at the University of Newcastle

People interested in Masters and PhD studies at the University of Newcastle can consult the "Research Higher Degree Candidate's Guide":
`http://www.newcastle.edu.au/research/rhd/guide/contents.html`
Information about funding and scholarships can be obtained from
`http://www.newcastle.edu.au/research/rhd/guide/schols.html`.
For enquiries about local scholarships or exchange arrangements please contact the school's office: School of Electrical Engineering and Computer Science, Faculty of Engineering and the Built Environment, The University of Newcastle NSW 2308, Australia Phone: +61 2 4921 5330, Fax: +61 2 4921 6929, URL: `http://www.eecs.newcastle.edu.au/` or send and email to the academic NUbot team members Dr. Stephan Chalup (`chalup@cs.newcastle.edu.au`) or Professor Richard Middleton (`rick@ee.newcastle.edu.au`).

The School of Electrical Engineering and Computer Science offers an attractive undergraduate programme which is well-suited to prepare for postgraduate research studies in the area of robotics. We list a small selection of the course descriptions of courses relevant for robotics. These descriptions are based on the 2003 course information linked the following page
`http://www.eng.newcastle.edu.au/eecs/current/courses.html`
where more undergraduate course information and details can be obtained.

- **Comp3330 Machine Intelligence.** This 3rd year elective course focuses on topics at the intersection of Artificial Intelligence and Machine Learning. Aim of the course is to provide a comprehensive introduction and overview of Machine Intelligence methods and algorithms with selected links to corresponding biological mechanisms. The course addresses the concepts of learning, searching, planning, interacting, reasoning, knowledge representation and intelligence. The emphasis of the course is on algorithmic aspects. But students will also learn about some software and applications. Central is the concept of an intelligent adaptive agent. We will discuss, for example, how agents can be designed to play strategy games or how to control soccer playing robots. The techniques of the course are very general. They are relevant for anyone who is interested in the development of future intelligent machines and programs. That could for example be robots exploring the surface of Mars or intelligent modules for business applications.
- **Comp4110 Advanced Machine Learning.** This is an Honours course for 4th year students. Through a series of advanced seminars and project work we will approach advanced topics from the areas of machine learning, data mining, neuroinformatics, and robotics.
- **Comp3320 Computer Graphics.** The beginning of the course focuses on selected fundamental geometric concepts which are broadly used in computer graphics. Then the lectures proceed with topics such as 3D-transforms, quaternions, optimisation, visual appearance, polygon manipulation, collision detection, special effects, lighting, colour, hypergraphics and computer

vision. In addition to the main stream of the lecture an overview about the history, tools and applications of computer graphics is acquired in a series of assignments. Practical experience with graphics programming is gained in the graphics and vision project.

– **Elec2120 Sensors and Actuators.** This course gives an introduction to the theory and practice of electrical and mechanical actuators and sensors. Examples to be studied are taken from the areas of electromagnetism, mechanical sensors, and piezoelectric and optical sensing systems. Elec2120 combines a theoretical background with practical experience of sensors, actuators and electronic transducers commonly used in measurement and control of modern industrial plants. Students are acquainted with the fundamentals of a wide range of common sensors and actuators, including an overview of DC and 3 phase AC motors as high power actuators. The course contains a significant component of practical laboratory and design experience.
– **Elec3710 Microprocessor Systems.** This course introduces students to the fundamental principles of designing embedded microprocessor systems. This involves learning all of assembly language programming on the Intel 80x86 architecture, 'C' language programming for embedded applications, handling interrupt driven I/O, the fundamentals of real time operating systems, and interfacing to I/O devices such as A/D and D/A converters.
– **Elec3850 Introduction to Electrical Engineering Design.** The devices developed in this course include a significant engineering component involving a range of disciplines including some or all of: Electrical, electronic, communications, computing, software, signal processing, control, and mechanical systems. Example products might include: (i) An exercise bike with regeneration into the AC mains; (ii) A simplified multiuser wireless telephone system; (iii) a "sumo" robot.
– **Elec4700 Advanced Computer Systems.** Introduces students to advanced concepts in computer architecture and design emphasing quantitative methods for performance evaluation. Topics include performance measures and cost, instruction set architecture, pipelining, instruction-level parallelism, caches, I/O and buses, interconnection networks.

Acknowledgements The NUbots are grateful to all colleagues, friends, previous members, and other supporters of the team. Many thanks to William McMahan and Jared Bunting who were temporary project students in Newcastle and contributed to the vision system. The NUbots received support from several sources, including: Research Infrastructure Grant; Faculty of Engineering & Built Environment Strategic Initiatives Fund; School of EE&CS Grant, and the Special Research Centre - The Centre for Integrated Dynamics and Control at the University of Newcastle in Australia.

Links to the NUbots' publications can be found at the NUbots' webpage

`http://robots.newcastle.edu.au/`

References

1. Jared Bunting, Stephan Chalup, Michaela Freeston, Will McMahan, Rick Middleton, Craig Murch, Michael Quinlan, Christopher Seysener, and Graham Shanks (2003). *Return of the NUbots ! The 2003 NUbots Team Report.* School of Electrical Engineering & Computer Science Technical Report, The University of Newcastle, Callaghan 2308, Australia.

2. Chalup, S., Creek, N., Freeston, L., Lovell, N., Marshall, J., Middleton, R., Murch, C., Quinlan, M., Shanks, G., Stanton, C., and Williams, M.-A. (2002). *When NUbots Attack ! The 2002 NUbots Team Report.* School of Electrical Engineering & Computer Science Technical Report, The University of Newcastle, Callaghan 2308, Australia.

3. Chernov, N. and Lesort, C. (2002). Least Squares Fitting of Circles and Lines, University of Alabama, Birmingham, USA. `http://arxiv.org/PS_cache/cs/pdf/0301/0301001.pdf`

4. Coleman, O.J. and Chalup, S.K. (2003). Towards Matching Perception in Simulated and Real World Robot Navigation. *Australian Conference on Artificial Life (ACAL).*

5. Freeston, L. (2002). *Localization and World Modelling.* `http://murray.newcastle.edu.au/users/students/2002/c9806118/main.html`

6. Freeston, M. (2003). *Localisation and World Modelling in Robot Soccer.* School of Electrical Engineering and Computer Science, The University of Newcastle, `http://murray.newcastle.edu.au/users/students/2003/c2003066/`

7. Hengst, B., Phan, S.B., Ibbotson, D., and Sammut, C. (2002). Omnidirectional Locomotion for Quadruped Robots. *RoboCup 2001: Robot Soccer World Cup V*, pp.368–373.

8. McMahan, W. and Bunting, J. (2002). *Vision Processing.* `http://murray.newcastle.edu.au/users/students/2002/c3012299/`

9. Quinlan, M., Murch, C., Middleton, R., and Chalup, S. (2003). Traction Monitoring for Collision Detection with Legged Robots. *RoboCup 2003 Symposium* (received engineering paper award).

10. Quinlan, M.J., Chalup, S.K., and Middleton, R.H. (2003). Techniques for Improving Vision and Locomotion on the Sony AIBO Robot. (submitted).

11. Quinlan, M.J., Chalup, S.K., and Middleton, R.H. (2003). Application of SVMs for Colour Classification and Collision Detection with AIBO Robots. *Neural Information Processing Systems (NIPS).*

12. Seysener, C. (2003). *Vision Processing for RoboCup 2003/2004.* School of Electrical Engineering and Computer Science, The University of Newcastle, `http://murray.newcastle.edu.au/users/students/2003/c9807674/index.html`