

UNSW United

UNSW-United

*James Westendorp, Malcolm Ryan, Philip Preston, Waleed Kadous,
John Dalgliesh, Mike Lawther, Graham Mann, Claude Sammut*

JW, MR, PP, WK, JD, ML, GM, CS: Artificial Intelligence Group, Computer Science & Engineering, University of New South Wales

Abstract. *The UNSW United Sony Legged robot team uses a number of techniques for learning, vision and planning. These include*

- *behavioural cloning;*
- *visual tracking of objects using a simple mechanism;*
- *a layered approach that allows alteration of local objectives;*
- *a high level visual interface to allow a human 'driver' to collect information for machine learning tools;*
- *augmentation of existing behaviours.*

1 Introduction

We use a variety of techniques for UNSW United's entry. The aim is to integrate the physical feature of the Sony Legged robots in an effort to allow flexible implementation of control algorithms based on learning and planning.

2 Behavioural Cloning

Given existing vision and localisation systems and the set of built in actions for walking, turning and kicking, more complex behaviours can be constructed using a hybrid of reinforcement learning and behavioural cloning. A human controller demonstrates both good and bad examples of the desired behaviour and the robot evaluates the controller's performance and thus learns the best way to produce that behaviour. Examples of behaviours that could be learned this way are passing, tackling and blocking. A higher level symbolic system can then combine these behaviours, along with other simpler ones programmed by hand, into plans of attack.

We make the assumption that the existing software for the Sony robots provides a starting point for the low-level behaviours, so that we can build on them. We are attempting to train the robots to perform high level behaviours. For example, we may set up a situation in which we want a robot to pass the ball to a team member. A human “trainer” will instruct the robot, be remote control operation, to perform the task. The trainer’s actions, plus the inputs to the robot, will be used for learning. In this case, the inputs would be the features extracted from the vision system, localisation information, etc. We have already used behavioural cloning in several domains, including learning to fly an aircraft in a flight simulator. This project are attempting a variant in which cloning is combined with reinforcement learning. In pure behavioural cloning, situation-action pairs are input to a program such as C4.5 to learn general situation-action rules. Here, a reinforcement learning algorithm uses these inputs to acquire its evaluation function.

3 Visual Control

Another variant of behavioural cloning explores the use of the robot’s vision system to create an iconic description of what it sees. These iconic descriptions are very compact, while simultaneously being very information-rich. This information can be used to give a “robot’s eye view” of the information available to it. This description constitutes a simple state representation. A human trainer then instructs the robot as to the appropriate action to take given this information. By repeating this process many times all the way to goal scoring, many situation-action pairs are generated. These can be fed into a learning algorithm to generate a classifier that can be used as a “clone” of the behaviour of the human trainer, by providing it with the iconic description. The classifier then recommends an action based on what it has learnt.

4 Vision

As vision is the major sensory modality for Sony legged robots in a soccer match, we have devoted substantial effort to the production of a range of support tools for calibration and image manipulation. These tools allow the discrimination of all the essential visual markers on the soccer field, but more importantly, allow calibration parameters to be rapidly calculated under a number of differing lighting conditions

5 Tracking

At a higher level, motion-tracking algorithms that we have been developing for other applications are also being applied. These have particularly been aimed for platforms with modest amounts of computational power. [Peters 1998] describes the basic techniques, which involve calculating a delta parameter from successive frames. These techniques are computationally very simple, and given the small image size of the Sony Legged Robot’s camera, this technique becomes a simple and effective tool.

6 Motion Control

We are also augmenting some of the basic robot behaviours, in order to facilitate more precise actions. The current range of motions is quite noisy, which is a problem when dealing with legged robots. Ideally, we would like some form of odometry in the robot. However, errors become very large very quickly, especially when applied over an extended period of time. However, 'local' odometry is an important form of measurement. Although we cannot eliminate noise from actions, we can aim to minimise it. More precise actions will enhance the reliability of the odometry, while simultaneously providing the robot with finer grained action sequencing. These new actions will also be more suited to the domain of robotic soccer.

7 Control Structure

At the highest level of abstraction, and perhaps the most important level, we are experimenting with a multi layered strategic planning approach incorporating elements from diverse disciplines of Artificial Intelligence. These include examples of expert systems and finite state machines in the upper levels of the hierarchy, with more traditional goal subdivision and sequencing closer to the effectors. Naturally consideration of the problems that these approaches entail is a serious issue. An example is the care that must be taken when applying heuristic knowledge so that there are no conflicts. This is an example of the classic expert system knowledge acquisition bottleneck. Techniques to counter problems of this nature are well known to the team and we have several useful solutions to hand [Compton et al 1993][Preston et al 1994].

Since team members in this league are unable to communicate in an effective fashion, cooperative behaviour, essential in any team activity, will have to be achieved by indirect means. This is a common problem in this domain [Werger et al 1997]. An example of this is perceiving a team member blocking an opponent's shot, and moving to a supporting role. The architecture we are adopting is amenable to this style of environmental interaction as it is based to some degree on Brooks' subsumption architecture [Brooks 86]. Given the concept of 'local objectives', 'remote objective' and 'global objectives', the overall task can be seen as a collection of immediately desired responses to the current situation. If teammates or opponents alter the current situation, then the current local objective can be altered

[1] Brooks, R (1986), A robust layered control system for a mobile robot, *IEEE Journal of Robotics and Automation*, V:RA-2:1, 14-23

[2] Compton, P., Kang, B., Preston, P. and Mulholland, M. (1993). Knowledge Acquisition without Analysis, in N. Aussenac, G. Boy, B. Gaines et al (Eds.), *Knowledge Acquisition for Knowledge Based Systems. Lecture Notes in AI (723)*. Berlin, Springer Verlag. 278-299.

[3] Mark W Peters and Arcot Sowmya WRAITH: Ringing the changes in a changing world, 1997, 4th Conference of the Australasian Cognitive Science Society, Newcastle, Australia.

[4] Preston, P., Edwards, G. and Compton, P. (1994). A 2000 Rule Expert System Without a Knowledge Engineer. *Proceedings of the 8th AAI-*

Sponsored Banff Knowledge Acquisition for Knowledge-Based Systems Workshop, Banff, Canada, 17.1-17.10.

[5] Sammut, C., Hurst, S., Kedzier, D. and Michie, D. (1992). Learning to Fly. In D. Sleeman & P. Edwards (Ed.), Proceedings of the Ninth International Conference on Machine Learning, Aberdeen: Morgan Kaufmann.