

Motivated Reinforcement Learning for Improved Head Actuation of Humanoid Robots

Jake Fountain, Josiah Walker, David Budden, Alexandre Mendes and
Stephan K. Chalup

The University of Newcastle, Australia
{jake.fountain, josiah.walker, david.budden}@uon.edu.au,
{alexandre.mendes, stephan.chalup}@newcastle.edu.au

Abstract. The ability of an autonomous agent to self-localise within its environment is critically dependent on its ability to make accurate observations of static, salient features. This notion has driven considerable research into the development and improvement of feature extraction and object recognition algorithms, both within RoboCup and the robotics community at large. Instead, this paper focuses on the rarely-considered issue imposed by the limited field of view of humanoid robots; namely, determining an optimal policy for actuating a robot’s head, to ensure it observes regions of the environment that will maximise the positional information provided. The complexity of this task is magnified by a number of common computational issues; specifically high dimensional state spaces and noisy environmental observations. This paper details the application of motivated reinforcement learning to partially overcome these issues, leading to an 11% improvement (relative to the null case of uniformly distributed actuation policies) in self-localisation and ball-localisation for an agent trained online for less than one hour. The method is demonstrated as a viable method for improving self-localisation in robotics, without the need for further optimisation of object recognition or tuning of probabilistic filters.

Keywords: motivated reinforcement learning, localisation, Fourier basis, head actuation, simulated curiosity

1 Introduction

Effective localisation is essential to solving tasks such as the Robocup soccer competition. The behaviour of a robot is a function strongly dependent on the robot’s localisation model. Inaccurate localisation leads to ineffective behaviour and poor performance. In a game of humanoid robot soccer, successful localisation of the robot and field objects relies on the vision system to measure relative field object locations. Moreover, the behaviour of the head determines the efficacy of the vision system; limited field of view implies that all field objects cannot be measured at once. The head actuation problem involves choosing neck motor positions to optimise localisation. If the robot is completely unlocalised,