Team Description: UW Huskies-03

C. T. Kwok, M. Burkhart, G. Hazen, M. Mohebbi, N. Stipanovich, and D. Fox

Department of Computer Science & Engineering
University of Washington, Seattle, WA

1 Introduction

This paper provides an overview of the robot control system of the UW Huskies, our entry in the RoboCup-2003 four-legged robot league. Our development team consists of one faculty member (D. Fox), one graduate student (C.T. Kwok, team leader), and five undergraduate students. As in the previous years, our focus is on robust and efficient state estimation using particle filters and computer vision techniques.

The control architecture follows a traditional, layered approach, shown in Figure 1. The different modules communicate using the Aperios inter-object message passing facility. When large data (e.g. world model, camera images) are to be shared between multiple modules, shared memory is used. The control system consists of the following five modules:

Hardware Interface: This module receives input from the raw sensors and passes it to the other modules. The most important information consists of time stamped camera images coupled with the corresponding configuration of the robot's head joints.

Vision: The vision module analyzes images on a frame by frame basis. Most low-level vision routines such as color segmentation and blob finding are based on the CMVision software package [2]. We additionally use color transitions to detect lines and field borders, which are used to improve the robot's position estimates (similar to [7]). The detected objects and lines are passed to the state estimator module.

State estimator: This module integrates information over time and stores it in the world model. Most important information is the position of the robot on the soccer field and the position of the ball relative to the robot. The robot position is estimated using a particle filter [5, 4]. Position estimates are based on the detection of markers, goals, lines, and borders of the field. A multi-model Kalman filter is used to estimate the relative position and velocity of the ball.

Behavior: Using the most recent information in the world model, this module decides which actions the robot should perform. Our control system follows a layered approach. Control behaviors are modeled by simple skills such as "kick" or "walk", and high-level decisions are made based on the current situation and the role of the robot. Roles are assigned dynamically based on the position of the robots on the field. Position estimates are generated by combining the estimates of all robots.

Motion: Low-level robot control uses the locomotion developed by the CMU CMPack team. Dead reckoning is based on experimentally determined velocity estimates.

Since most of the techniques are already described in previous reports [3], we will focus on our novel approach to ball tracking.

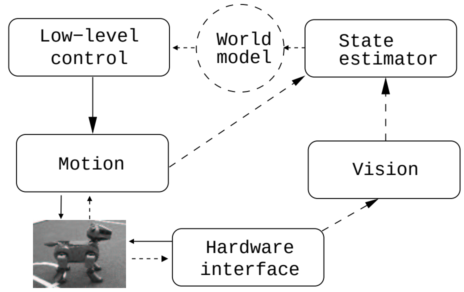

Fig. 1. System architecture. The arrows indicate the main flow of information (dashed) and commands (solid) between the different modules.

2 Multi-model Ball Tracking

Due to their efficiency Kalman filters are extremely well suited for state estimation in RoboCup. This is especially true for tracking the location of the ball since the uncertainty in the ball's position is typically unimodal. Unfortunately, the dynamics of the ball are highly non-linear, which makes the application of Kalman filters not straightforward. For example, if the ball is in a free motion on the field, its trajectory is almost linear. If, however, the ball gets kicked by another robot or if it is grabbed, the motion pattern of the ball is very different from free motion. We address this problem by concurrently estimating the ball's position and velocity *and* the discrete mode of motion (rolling, not moving, kicked, grabbed). Such a joint estimate over mixed discrete and continuous spaces can be maintained efficiently using Rao-Blackwellised particle filters [6, 1, 8]. Here, the discrete states of the ball are generated using the sampling, importance sampling with re-sampling procedure of particle filters [4]. Each particle is annotated with a Kalman filter which estimates the position of the ball *conditioned on the mode of the corresponding particle*. In our first experiments this approach yields highly superior performance over standard Kalman filters while still being efficient enough to run onboard the robot.

3 Conclusions and Ongoing Work

Our software system for RoboCup-2003 relies on well-established techniques for robot control and state estimation. Currently, we investigate a sample based implementation of active sensing. The key idea of our approach is to determine the robot's sensing direction by minimizing the uncertainty in the quantities that are most relevant in the current situation. For example, if the robot is far away from the ball, an accurate estimate of the ball's position is not needed. If, however, the robot wants to kick the ball into the goal, it needs to accurately know the location of the ball and the direction to the goal (not necessarily the distance). We intend to use reinforcement learning to determine the importance of different projections of the robot's state space. Another area of research is unsupervised color labelling. Similar to the joint estimation approach for multi-model ball tracking, we concurrently estimate the position of the robot's camera on the field *and* the labelling for the color values.

Acknowledgments

This research is sponsored in part by NSF (CAREER contract number 0093406).

References

1. C. Andrieu, M. Davy, and A. Doucet. Efficient particle filtering for jump markov systems. In *IEEE International Conference on Acoustics Speech and Signal Processing*, 2002.
2. J. Bruce, T. Balch, and M. Veloso. Fast and inexpensive color image segmentation for interactive robots. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2000.
3. Z. Crisman, E. Curre, C.T. Kwok, L. Meyers, N. Ratliff, L. Tsybert, and D. Fox. Team description: UW Huskies-02. In G. Kaminka, P.U. Lima, and R. Rojas, editors, *RoboCup-2002: Robot Soccer World Cup VI*. Springer Verlag.
4. D. Fox. Adapting the sample size in particle filters through KLD-sampling. *International Journal of Robotics Research (IJRR)*, 22, 2003.
5. J.S. Gutmann and D. Fox. An experimental comparison of localization methods continued. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2002.
6. K. Murphy and S. Russell. Rao-Blackwellised particle filtering for dynamic Bayesian networks. In A. Doucet, N. de Freitas, and N. Gordon, editors, *Sequential Monte Carlo in Practice*, New York, 2001. Springer-Verlag.
7. T. Röfer and M. Jüngel. Fast and robust edge-based localization in the sony four-legged robot league. In *7th International Workshop on RoboCup 2003*, Lecture Notes in AI. Springer, 2003.
8. D. Schulz, D. Fox, and J. Hightower. People tracking with anonymous and id-sensors using Rao-Blackwellised particle filters. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2003.