Policy Learning for Autonomous Feature Tracking

Maria Fox

Department of Computer and Information Sciences,
University of Strathclyde, Glasgow, G1 1XH, UK
email: maria.fox@cis.strath.ac.uk

Invited Talk
3rd ICAPS Workshop on Validation and Verification of Planning and Scheduling Systems (VVPS 2011)
Freiburg, Germany, June 13, 2011

When operating under uncertainty, a robust behaviour can often be learned by sampling a large number of cases, planning what to do precisely in those cases (which are deterministic) and then learning a general policy for what to do in future unseen cases. The learning can be done by classifying the actions that can be performed into the states that were met in the training cases. The resulting policy will be good if the sampling was representative.

This talk presents an example of this approach applied to the problem of tracking a moving patch of water in the ocean. The same approach has been successfully applied to the problem of load management for multiple independent batteries. The talk will address the modelling problems that arise in the patch-tracking case, some of the subtleties that affect the quality of the learned policy, and the performance of the system executing the learned policy. The results we have obtained so far suggest that the planning followed by classification approach is highly scalable and is a practical and effective method for achieving intelligent behaviour under uncertainty.