Feature selection for domain knowledge representation through multitask learning

Benjamin Rosman
Mobile Intelligent Autonomous Systems
CSIR
South Africa
BRosman@csir.co.za

Representation learning is a difficult and important problem for autonomous agents. This paper presents an approach to automatic feature selection for a long-lived learning agent, which tackles the trade-off between a sparse feature set which cannot represent stimuli of interest, and rich feature sets which increase the dimensionality of the space and thus the difficulty of the learning problem. We focus on a multitask reinforcement learning setting, where the agent is learning domain knowledge in the form of behavioural invariances as action distributions which are independent of task specifications. Examining the change in entropy that occurs in these distributions after marginalising features provides an indicator of the importance of each feature. Interleaving this with policy learning yields an algorithm for automatically selecting features during online operation. We present experimental results in a simulated mobile manipulation environment which demonstrates the benefit of our approach.