Models of Behavioural Learning

Organisers:	Chris Watkins	Royal Holloway, UK
	Nicolo Cesa-Bianchi	University of Milan, Italy
	Stefan Schaal	USC
	John Shawe-Taylor	University of Southampton, UK

Sponsored by the Pascal Network of Excellence.

Aims

The main aim of the workshop is to raise basic questions about reinforcement learning (RL). The format will be longer contributed and tutorial talks, punctuated by considerable discussion. A secondary aim is to provide tutorial surveys of as wide a range of models of behavioural learning as possible: animal learning phenomena; RL models of neural learning; non-stochastic RL for bandit problems; motor learning; and new approaches to learning in POMDPs.

The morning session will concentrate on RL and animal learning.

The notion that animals learn to obtain rewards and to avoid punishments might seem so natural and obvious that it is hardly worth discussing. Actually, for much — perhaps most — animal learning, it may not even be true.

There are well known examples of animal learning that are difficult to interpret as RL: are these examples anomalies? Or is RL itself only a small part of a larger explanatory theory of animal learning that has not yet been properly articulated? Should machine learning researchers worry that so much of animal learning does not appear to be naturally formulated as RL?

In the morning, Chris Watkins will give a brief introduction and raise some questions. Nathaniel Daw (Gatsby Unit, UCL, UK)will give a critical survey of application of RL models to animal learning and neurobiological data. The next talk is given by Randy Gallistel (Rutgers), who is the author of seven books and over 100 journal articles mostly on animal learning. He will give a tutorial survey of animal learning phenomena, including those difficult to interpret in RL terms. Nicolo Cesa-Bianchi will then give a tutorial

on non-stochastic approaches to bandit learning problems, for which there are simple algorithms that might seem ideally suited for animal use.

In the afternoon, the focus will be on learning and robotics. Yishay Mansour (Tel Aviv) will describe recent work on learning and planning in POMDPs. Stefan Schaal will give a tutorial survey of models of motor learning in robots. Susanne Still (Hawaii) will describe a novel approach to behavioural learning based on the notion that a learner should maximise the predictive information in its internal model. Finally, Peter Auer will describe an approach to the competitive learning of policies, which is a new departure in RL.

Morning Session

7:30-8:00	Chris Watkins	Welcome and Overview: Questions about RL
8:00-8:40	Nathaniel Dawe	RL in the Brain: models and evidence + discussion
8:40 - 9:40	Randy Gallistel	A Survey of Animal Learning Phenomena + discussion
9:40 - 10:30	Nicolo Cesa-Bianchi	Competitive Learning in Bandit Problems.

Afternoon Session

3:30-4:15	Yishay Mansour	Learning and Planning for POMDPs
4:15-5:00	Stefan Schaal	Learning Motor Behaviors — Past and Present Work
5:00-5:30	Susanne Still	Active Learning and Optimal Predictions
5:30-5:50	Peter Auer	Competitive Learning of Policies
5:50-6:30	General discussion and wrap-up	

If you would like to make a contribution to this workshop, plese contact Chris Watkins `chrisw@cs.rhul.ac.uk`