Centre for Cognition, Computation and Modelling, Birkbeck
One-day Workshop:

Associative Learning, Connectionism & Reinforcement Learning: Past, Present, Future

Monday June 1st[th], Birkbeck, University of London
Time: 10:00
Place: room B33 Malet Street, Birkbeck Main Building, Torrington Square

Attendance free, though we would appreciate indication of intent to attend to: u.hahn@bbk.ac.uk

Speakers:

John Pearce, Cardiff University

Jay McClelland, Stanford University

Peter Dayan, Gatsby Computational Neuroscience Unit

Koray Kavukcuoglu, Google DeepMind

Denis Mareschal, Birkbeck

Abstracts

John Pearce:

Associative Learning: Past, Present and Future

The study of associative learning in animals has been dominated by two questions. What knowledge is acquired during an associative learning episode, and what are the circumstances that promote the acquisition of this knowledge. As far as the second question is concerned, studies of Pavlovian conditioning have shown that associative learning will take place when a neutral stimulus is followed by a biologically significant outcome whose occurrence is unexpected or surprising. The surprisingness of the outcome is determined by how well it is predicted to occur on the basis of all the stimuli that are present at the time of its delivery. A further factor that influences the effectiveness of associative learning is the amount of attention paid to stimuli that signal the delivery of the important outcome. Attention has been shown to vary according to whether the

stimuli have been reliably associated with an important outcome in the past, and according to whether the stimuli have recently been followed by unexpected events. As for the knowledge that is acquired during Pavlovian conditioning, it is said to involve associations between representations of the neutral stimulus and the biologically important event, so that the activation of the former is able to activate the latter, and thus indirectly excite responses appropriate to the imminent delivery of the outcome. It will be argued that the representation of the signal for the biologically important event is based on the entire pattern of stimulation that immediately precedes its delivery.

Jay McClelland:

Many learning mechanisms: Deep neural networks and beyond

Human behavior is influenced by prior knowledge at very short latencies and often outside of cognitive control, suggesting that one of the consequences of experience is the structuring of wired-in tendencies to map inputs to internal representations and responses, similar to the process of learning connections in multilayer neural networks used to perform such tasks as object and speech recognition and/or rapid responding under time pressure in real-world, computer game, or behavioral experimental settings. I have argued that such processes also shape our comprehension of linguistic and mathematical symbols, and more generally the representations and reactions that become active in our minds when we perceive or remember, and even form the basis for our firmly held beliefs and intuitions. These tendencies are acquired gradually, however, over developmental time, and there must be more to learning than this. Sometimes we can learn from only one or a few experiences. Furthermore, in situations were we have time to think and consider relationships between experiences, connection adjustment is unlikely to be the only mechanism whereby we distill knowledge from experience. I will discuss these points and then describe some experiments from my laboratory in which we found evidence, under different task conditions, of *three* different ways in which experience can produce learning that then guides responses to familiar and novel situations. In one of these cases, I find it plausible to suppose participants employ explicit reasoning to produce inferences which then can be stored in memory to guide future behavior.

Peter Dayan:

Statistical Models of Conditioning

Animal and human conditioning has long been a very attractive target for computational models. The most venerable models are forms of learning rule, endowed with sophisticated features that have been designed in the light of some of the wealth of empirical observations. A more recent trend has been to derive learning rules from a variety of statistical principles. I will talk about some of our and others' work in the latter direction, including ideas about associability, stimulus competition and contexts.

Koray Kavukcuoglu:

Deep Reinforcement Learning

Reinforcement learning agents have achieved some successes in a variety of domains, however their applicability has previously been limited to domains in which useful features can be handcrafted, or to domains with fully observed, low-dimensional state spaces. In this talk I will explain two novel approaches that combine deep learning and reinforcement learning to enable agents to derive efficient representations of the environment from high-dimensional sensory inputs, and use these to generalize past experience to new situations. The Deep Q-Network (DQN) algorithm that achieves human level performance on ATARI 2600 domain and a Recurrent Attention Model (RAM) that learns a policy for visual glimpses to solve sequence prediction problems directly from images.

Denis Mareschal: Discussant

Full Schedule:

10:00 - 11:00 John Pearce
11:00 - 11:15 Break
11:15 - 12:15 Jay McClelland

12:15 – 13:30 Lunch Break

13:30 – 14:30 Peter Dayan
14:30 - 14:45 Break
14:45 - 15:45 Koray Kavukcuoglu
15:45 - 16:30 Denis Mareschal & Open Discussion