BCL5

5th Workshop on Brain, Computation, and Learning

Maneesh Sahani

Professor of Theoretical Neuroscience and Machine Learning
Gatsby Computational Neuroscience Unit
University College London

1st July | 14:00 – 15:30
Biological Sciences Auditorium

Towards a more biologically plausible connectionism

Abstract: We are living in the heyday of connectionist machine learning. Neural-network-based machines, with simple homogeneous computational architectures encompassing billions of parameters and trained using enormous data sets, have proven to be amazingly capable. At times they appear to emulate the knowledge and even reasoning of an expert human being. Does this mean they are good models of biological intelligence? I will argue not. Some of the challenge lies in the lack of external referents during training (and may be fixed relatively soon). More substantively, however, neither the architecture, nor the learning rules, nor the training data, nor the objective functions are biologically realistic. Thus, even if the AI systems' reasoning can be better grounded in reality, quite fundamental gaps to biological processes are likely to remain.

Can this be fixed? I will sketch our attempts to do so. The goal is to understand an alternative type of inferential connectionist architecture, that uses local learning rules to build accurate models of the world from sensory (and sensorimotor) experience, and uses recurrently coupled distributed dynamical systems to integrate noisy and ambiguous sensory input with such internal models. I will focus on mathematical elements of the new theory, highlighting the key results that enable biologically plausible inference, learning and structure discovery.