



McGill McGill Seminar in Speech and Language Technology

Speaker: Roger Moore

Affiliation: Professor, Department of Computer Science, University of Sheffield

Date: Monday, June 21, 2010

Time: 11:00 AM

Location: Room 603, McConnell Engineering Building, 3480 University Street

Spoken Language Processing: Where Do We Go From Here?

Recent years have seen steady improvements in the quality and performance of speech-based human-machine interaction driven by a significant convergence in the methods and techniques employed. However, the quantity of training data required to improve state-of-the-art systems seems to be growing exponentially, yet performance appears to be reaching an asymptote that is not only well short of human performance, but which may also be inadequate for many real-world applications. This suggests that there may be a fundamental flaw in the underlying architecture of contemporary systems, and the future direction for research into spoken language processing is currently uncertain.

This talk addresses these issues by stepping outside the usual domains of speech science and technology, and instead draws inspiration from recent findings in the neurobiology of living systems. In particular, four areas will be discussed: the growing evidence for an intimate relationship between sensor and motor behaviour in living organisms, the power of negative feedback control to accommodate unpredictable disturbances in real-world environments, mechanisms for imitation and mental imagery for learning and modeling, and hierarchical models of temporal memory for predicting future behaviour and anticipating the outcome of events.

The talk will conclude by showing how these results point towards a novel architecture for speech-based human-machine interaction that blurs the distinction between the core components of a traditional spoken language dialogue system; an architecture in which cooperative and communicative behaviour emerges as a by-product of a model of interaction where the system has in mind the needs and intentions of a user, and a user has in mind the needs and intentions of the system.

Roger Moore has been professor of spoken language processing in the Speech and Hearing Research Group at the University of Sheffield since 2004 and is a visiting professor in the Department of Linguistics and Phonetics at University College London. Previously, he was chief scientific officer at 20/20 Speech and before that Head of the UK government Speech Research Unit.

He is also past president of both the International Speech Communication Association and the Permanent Council of the International Conferences on Spoken Language Processing. He is currently the editor of the journal of Computer Speech and Language.