

Keynote: Statistical Approaches to Open-domain Spoken Dialogue Systems

Steve Young
Cambridge University, UK
sjy@eng.cam.ac.uk

In contrast to traditional rule-based approaches to building spoken dialogue systems, recent research has shown that it is possible to implement all of the required functionality using statistical models trained using a combination of supervised learning and reinforcement learning. This approach to spoken dialogue is based on the mathematics of partially observable Markov decision processes (POMDPs) in which user inputs are treated as observations of some underlying belief state, and system responses are determined by a policy which maps belief states into actions.

Virtually all current spoken dialogue systems are designed to operate in either a specific carefully defined domain such as restaurant information and appointment booking, or they have very limited conversational ability such as in Siri and Google Now. However, if voice is to become a significant input modality for accessing web-based information and services, then techniques will be needed to enable conversational spoken dialogue systems to operate within open domains.

This talk will discuss methods by which current statistical approaches to spoken dialogue can be extended to cover much wider domains. It will be argued that unlike many other areas of machine learning, spoken dialogue systems always have a user on-hand to provide supervision. Hence spoken dialogue systems provide a unique opportunity to automatically adapt on large quantities of speech data without the need for costly annotation.