CoE-MaSS weekly seminar series

THE DST-NRF CENTRE OF EXCELLENCE IN MATHEMATICS AND STATISTICAL SCIENCES (CoE-MaSS) WOULD LIKE TO PRESENT A RESEARCH SEMINAR BY

Dr Andrew Saxe

(Center for Mind, Brain and Computation, Stanford University, and Visiting Researcher at Wits University)

"Demystifying deep learning: Learning dynamics in deep linear neural networks"

Friday, 19 June 2015
10h30-11h30
Videoconferencing Facility, First Floor
Mathematical Sciences Building, Wits West Campus

Tea will be served in the MSB Staff Room on the Upper Ground Floor from 09h45-10h30 for those that are attending the seminar at Wits.

How to connect to this seminar remotely:
You can connect remotely via Vidyo to this research seminar by clicking on this link:
http://wits-vc.tenet.ac.za/flex.html?roomdirect.html&key=y0SSOwFsvsidbzg4qFdWXvvQtyI
and downloading the Vidyo software before the seminar.You can join in the virtual venue (called "CAM Seminar Room" on Vidyo) to check your settings beforehand, from 10h00-10h15.
Important videoconferencing netiquette: Once the seminar commences, please mute your own microphone so that there is no feedback from your side into the virtual room. During the Q&A slot you can then unmute your microphone if you have a question to ask the speaker.

Title:

Demystifying deep learning: Learning dynamics in deep linear neural networks

Abstract:

Humans and other organisms show an incredibly sophisticated ability to learn about their environments during their lifetimes. This learning is thought to alter the strength of connections between neurons in the brain, but we still do not understand the principles linking synaptic changes at the neural level to behavioral changes at the psychological level. Part of the difficulty stems from depth: the brain has a deep, many-layered structure that substantially complicates the learning process.

To understand the specific impact of depth, I develop the theory of gradient descent learning in deep linear neural networks. Despite their linearity, the learning problem in these networks remains nonconvex and exhibits rich nonlinear learning dynamics. I give new exact solutions to the dynamics that quantitatively answer fundamental theoretical questions such as how learning speed scales with depth. These solutions revise the basic conceptual picture underlying deep learning systems--both engineered and biological--with ramifications for a variety of phenomena.

As time permits, I will then highlight three consequences at different levels of detail. First, the theory shows that layerwise unsupervised learning is a domain general strategy for speeding up subsequent learning, which I link to critical period plasticity in sensory cortices. Second, the theory suggests that depth influences the size and timing of receptive field changes in visual perceptual learning. And third, by considering data drawn from structured probabilistic graphical models, the theory reveals that only deep (and not shallow) networks undergo quasi stage-like transitions during learning reminiscent of those found in human semantic development. These applications span levels of analysis from single neurons to cognitive psychology, demonstrating the potential of deep linear networks to connect detailed changes in neuronal networks to changes in high-level behavior and cognition.