Prof. Wulfram Gerstner
EPFL – IC & SV– LCN
SG –AAB119 (Bâtiment SG-AAB)
Station 15
CH - 1015 LAUSANNE, Switzerland

Thursday, May 10ᵗʰ, 2012

Swiss Computational Neuroscience Series in ETHZ, Zurich

13h00 - Prof. Kenneth HARRIS
Imperial College London

The neural marketplace

This talk will describe a hypothesis for neuronal self-organization, in which competition for retroaxonal factors causes neurons to form functional networks through processes similar to those of a free-market economy.

Classically, neurons communicate by anterograde conduction of action potentials. However, information can also pass backward along axons, a process that is well characterized during the development of the nervous system. Recent experiments have shown that information about changes to a neuron's output synapses may pass backward along the axon, and cause changes in the same neurons inputs. Here we suggest a computational role for such "retroaxonal" signals in adult learning. We hypothesize that strengthening of a neuron's output synapses stabilizes recent changes in the same neuron's inputs. During learning, the input synapses of many neurons undergo transient changes, resulting in altered spiking activity. If this in turn promotes strengthening of output synapses, the recent synaptic changes will be stabilized; otherwise they will decay. A representation of sensory stimuli therefore evolves that is tailored to the demands of behavioral tasks.

The talk will describe a mathematical model of this process in recurrent networks of excitatory and inhibitory sigmoidal firing rate neurons. We define an objective function consisting of an output error plus a penalty term measuring how much the network's output would be altered by a random perturbation of its inputs. The penalty term has the effect of enforcing output smoothness, and in the case of a linear feedforward network, reduces to ridge regression. We derive an approximation for the penalty term in terms of flows of retroaxonal "payments" proportional to synaptic weights. Gradient ascent of the objective function causes hidden neurons to perform a form of unsupervised learning by which they avoid intermediate values of input current that render them more vulnerable to perturbations in their inputs. We consider a second, non-gradient step, by which neurons sporadically evaluate large changes in their input weight distributions. We show that optimization of this step leads to changes being retained if it causes an increase in the neuron's "profit," which consists of its retroaxonal payments minus those it must pay to its own inputs. Using simulations, we confirm the validity of the approximations made by the mathematical theory.

Finally, we describe the beginnings of an implementation of this scheme in spiking neurons. Specifically, we derive an unsupervised learning rule we term the "convallis rule" based on similarities to that derived above, and show it allows a subsequent output layer to produce good performance on a real-world speech recognition task.