In week 6 of FNCM'13 you will work in groups on a small project studying an extension of one of the models we looked at in the lectures and computer labs. The projects will be carried out in a group, and presented in the sessions of Friday Nov. 29th and Monday Dec. 2nd.

The proposed projects (with thanks to Phong) are the following - you may in your group adapt the project based on your interests and the technical possibilities & capabilities.

1. Hodgkin-Huxley. In the computer lab we only looked at the 2-dimensional Fitzhugh-Nagumo model, but the R-package we used allows for simulating the dynamics of more dimensional dynamical systems as well, so you may study the dynamics of the original 4-dimensional Hodgkin-Huxley model to evaluate the appropriateness of the simplification in the FN-model.

2. Fitzhugh-Nagumo. The HH- and FN-model greatly simplify the description of a neuron by ignoring the spatial extension of the neuron. In compartment models you can put some of that spatial extension back in by describing a neuron as a set of coupled compartment, treating the electrical activity of one compartment as input current to its neighboring compartments. This way you can model e.g., a running wave along the axon.

3. Multilayer perceptron. An interesting variant of the MLP is an autoencoder, where the target activation of the output units is set such that it equals the input. Behavior of such networks can be compared with principle component analysis (available in R) and there are interesting connections with structures in real brains (in the human brain, layers of neurons are stacked in order to extract abstract, semantic features from noisy and redundant input signals. For instance, in the visual system, neurons in V1 are for edge and corner detection whereas neurons in higher layers selectively respond to complex objects.)

Recommended reading:
- UFLDF tutorial: http://ufldl.stanford.edu/wiki/index.php/UFLDL_Tutorial
- Hinton, Geoffrey E., and Ruslan R. Salakhutdinov. "Reducing the dimensionality of data with neural networks." Science 313.5786 (2006): 504-507.

4. Structure representation
How neural networks store and process symbolic structures is still an open question. In his talk, Gideon presented dynamic binding, especially his work on hierarchical prediction network, as a solution. In this project, you may look at the Recursive Auto-Associative Memory (RAAM) of Pollack, a recurrent model that shares some features with the autoencoder of project 3.

Recommended reading:
- Pollack, Jordan B. "Recursive distributed representations." Artificial Intelligence 46.1 (1990): 77-105.
- Section 4.3 in Gideon's PhD thesis: The Neural Basis of Structure in Language: bridging the gap between symbolic and connectionist models of language processing (ILLC Dissertation Series DS-2011-11)

5. Spurious patterns in Hopfield net.
In the week 4 computer lab, we see that a Hopfield net sometimes converges to some unexpected patterns, which are called spurious patterns. In this project, you will study why those patterns occur and how to remove them.

Recommended reading:
- Hopfield, John J., D. I. Feinstein, and R. G. Palmer. "‘Unlearning’ has a stabilizing effect in collective memories." (1983): 158-159.

6. Towers of Hanoi. In week 5 computer lab we saw a simple implementation of an algorithm that computes the optimal solution for a Towers of Hanoi problem, with the restriction that all 4 rings are on a single peg both in the initial and in the target configuration (so called tower-to-tower problems). In this project you remove this restriction so that you can deal with flat-to-tower, tower-to-flat and flat-to-flat problems, and report on the size and contents of the stack during the computation and execution of a plan. Can you assess whether the Altmann & Trafton model will make as good predictions for such problems?

A final project I mention here for inspiration, but it is probably too ambitious to connect this to a computer implementation in the few days available.

7. Is backpropagation biologically plausible?
Backpropagation was certainly a breakthrough in connectionsm, but many received wisdom is that the backward propagation of error is biologically implausible. Recently, however, this received wisdom has been challenged. The purpose of this project is to investigate the question "Is backpropagation biologically plausible?".

Recommended reading:
- Section 2.3 in Thomas, Michael SC, and James L. McClelland. "Connectionist models of cognition." The Cambridge handbook of computational psychology (2008): 23-58.
- O'Reilly, Randall C. "Six principles for biologically based computational models of cortical cognition." Trends in cognitive sciences 2.11 (1998): 455-462.
- Xie, Xiaohui, and H. Sebastian Seung. "Equivalence of backpropagation and contrastive Hebbian learning in a layered network." Neural Computation 15.2 (2003): 441-454.