Models of Olfactory Information Processing

Dr. Thomas Nowotny
University of Sussex

Why are we interested in olfaction of insects?

- Biological sensory systems have an amazing performance
- Sensory systems are more accessible than deeper brain regions (controlled input space)
- Insect olfaction is a good model to understand sensory processing
  - The systems are comparably small and experimentally accessible
  - Structure is very similar across species
  - Many recent advances (Nobel prize 2004)

Main olfactory pathway anatomy

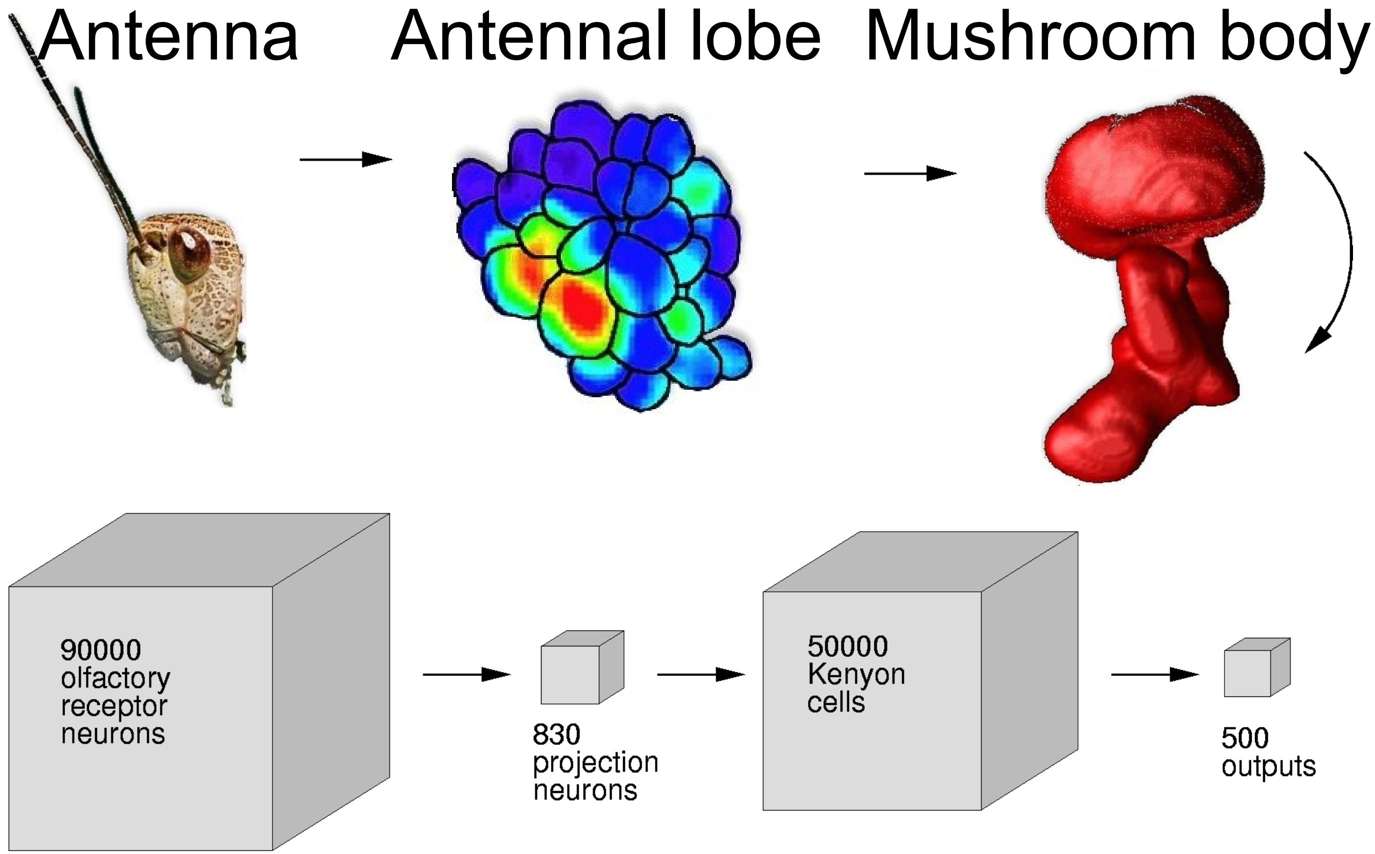

Box volume ~ number of cells

Dr. Thomas Nowotny,
Centre for Computational Neuroscience and Robotics
University of Sussex

Olfactory Receptors

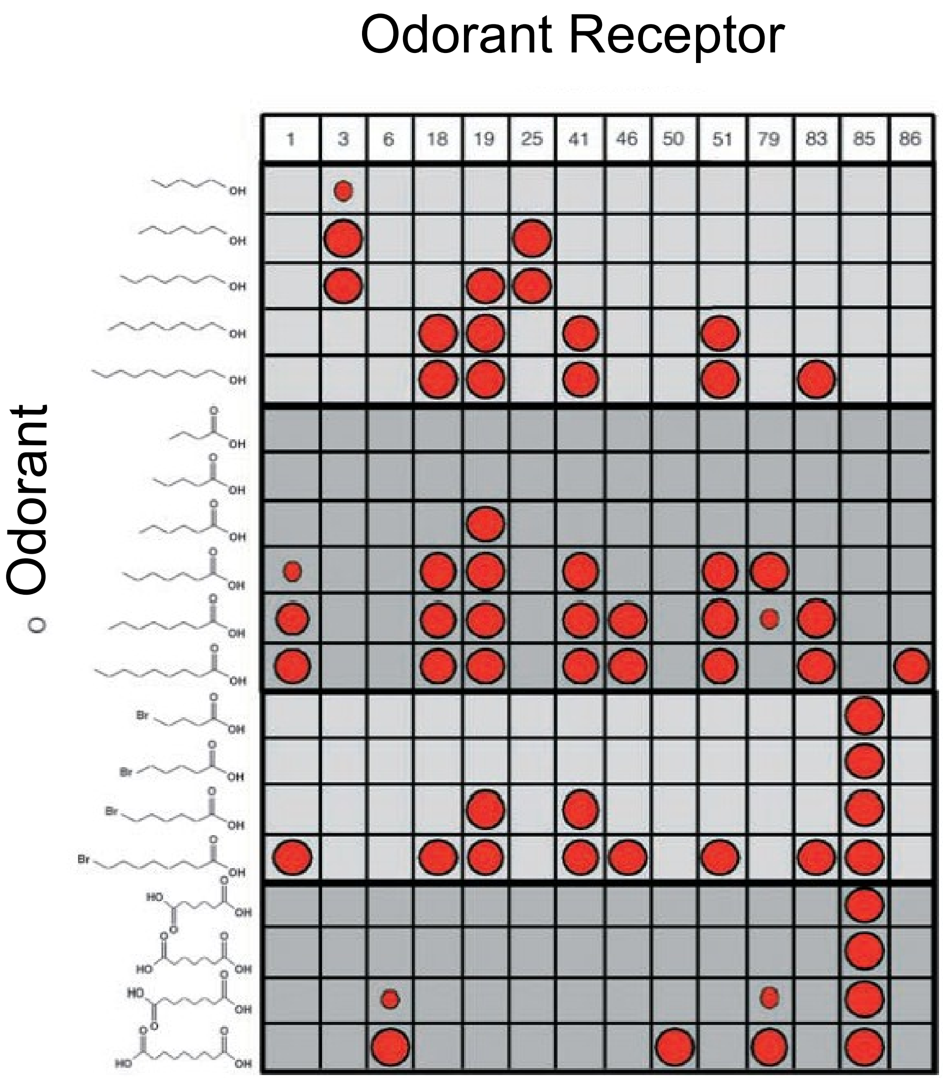

Odors evoke different, but overlapping patterns of receptor activity

From Linda Buck: Nobel lecture

Early processing

- Each olfactory receptor neuron expresses **one** receptor type
- All olfactory receptor neurons of the same type converge onto **the same** glomerulus
- Projection neurons receive inputs from **one** glomerulus

Odors are encoded as overlapping patterns of projection neuron activity.

In Richard Axel's words

“The elucidation of an olfactory map […] leaves us with a different order of problems. Though we may look at these odor-evoked images with our brains and recognize a spatial pattern as unique and can readily associate the pattern with a particular stimulus, the brain does not have eyes. “

Richard Axel, Nobel lecture

In other words:

The algorithm of olfactory information processing remains to be found.

Aside: PN patterns are processed in “snapshots”

There is a complex spatial-temporal dynamics in the AL, **but**

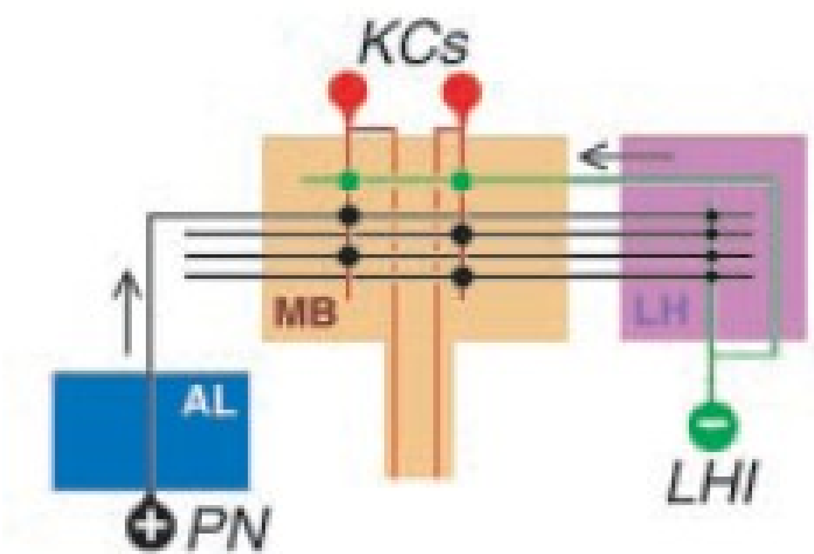

Local Field Potential corresponds to a periodic 20 Hz inhibition onto KCs in the MB

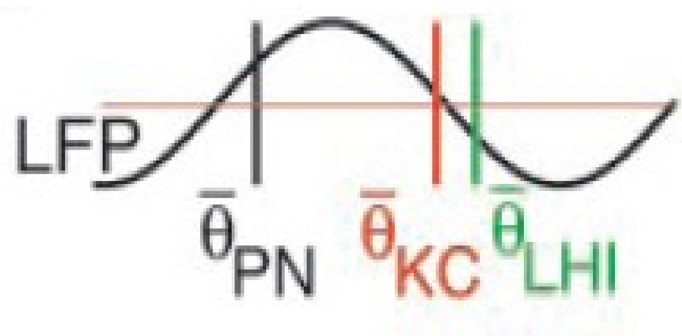

Perez-Orive et al., Science (2002)

We will assume in the following that odor information is transformed into time discrete “snapshots” of activity patterns transmitted to the KC of the mushroom body.

A classical pattern recognition solution: Perceptron

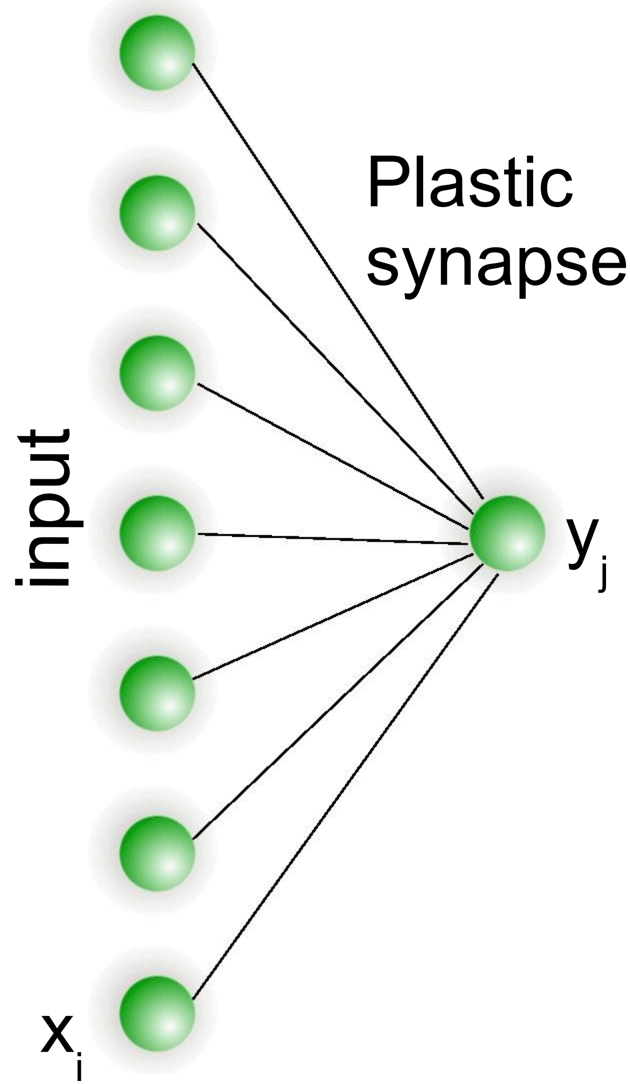

A simple perceptron rule:

Train y to respond to odor **x**
(call it class 1)

… and hope that y does not respond to *any other odor*
(call it class -1)

Typical implemetation:
McCulloch-Pitts neurons

$$y_j(t) = \Theta\big(\sum_i w_{ji} x_i(t-1) - \theta\big)$$

McCulloch-Pitts neurons are hyperplanes

The equation $\sum_{j=1}^{N} w_{ij} x_j = \theta$ defines a plane in N dimensional space.

For example N=2:

$$w_{i1} x_1 + w_{i2} x_2 = \theta$$

$$\Leftrightarrow x_2 = -\frac{w_{i1}}{w_{i2}} x_1 + \frac{\theta}{w_{i2}}$$

$$\Leftrightarrow y = ax + b \qquad \left(a = -\frac{w_{i1}}{w_{i2}},\ b = \frac{\theta}{w_{i2}} \right)$$

“McCulloch-Pitts neurons fire to the right of a hyper-plane and are silent on the left.”

This perceptron is a linear classifier

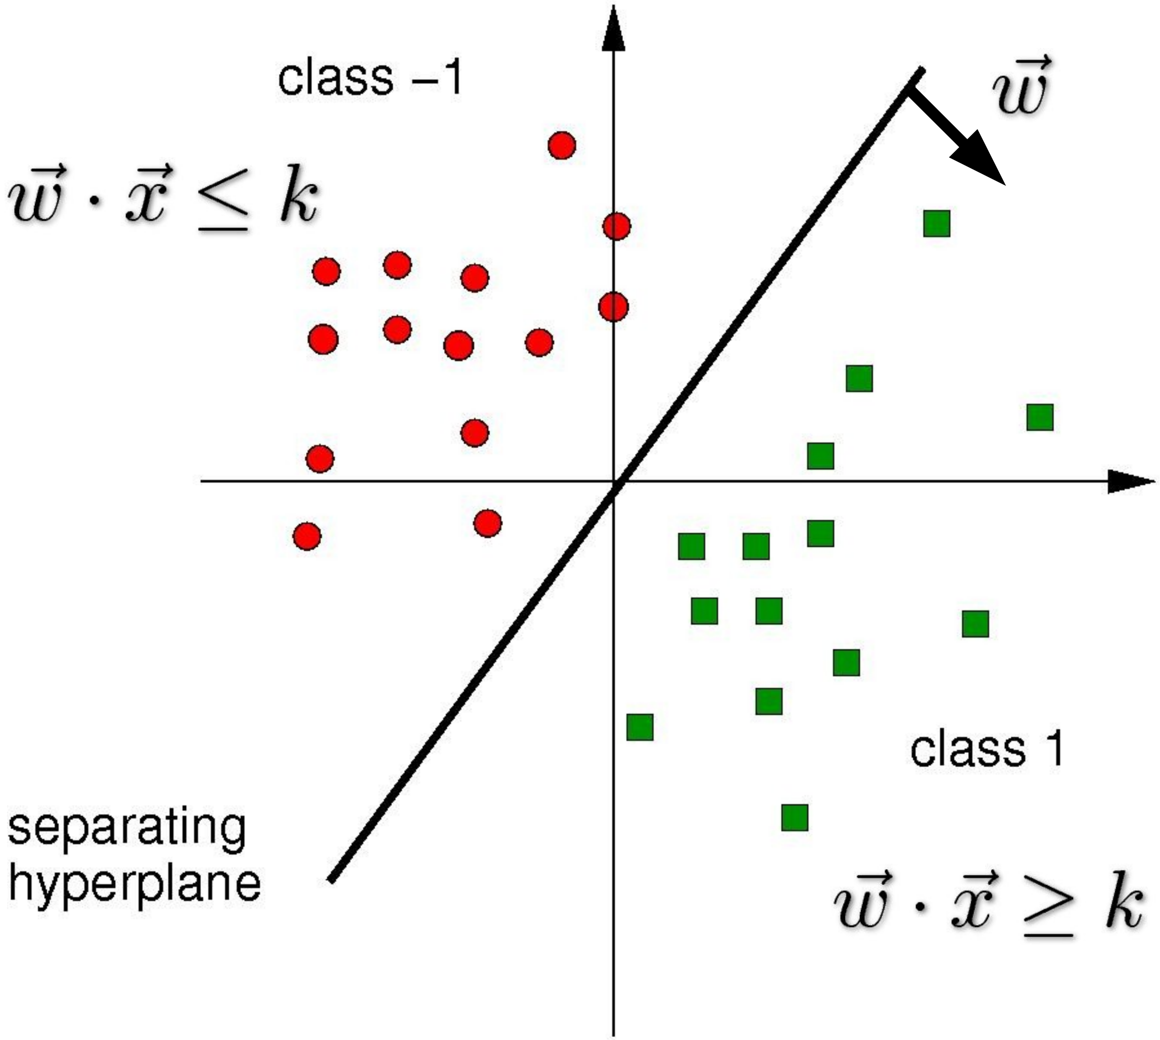

The hyperplane is adjusted through the training and Hebbian learning

Dr. Thomas Nowotny,
Centre for Computational Neuroscience and Robotics
University of Sussex

Support Vector Machines (SVM)

Cortes and Vapnik 1992,95: Support vector machine:

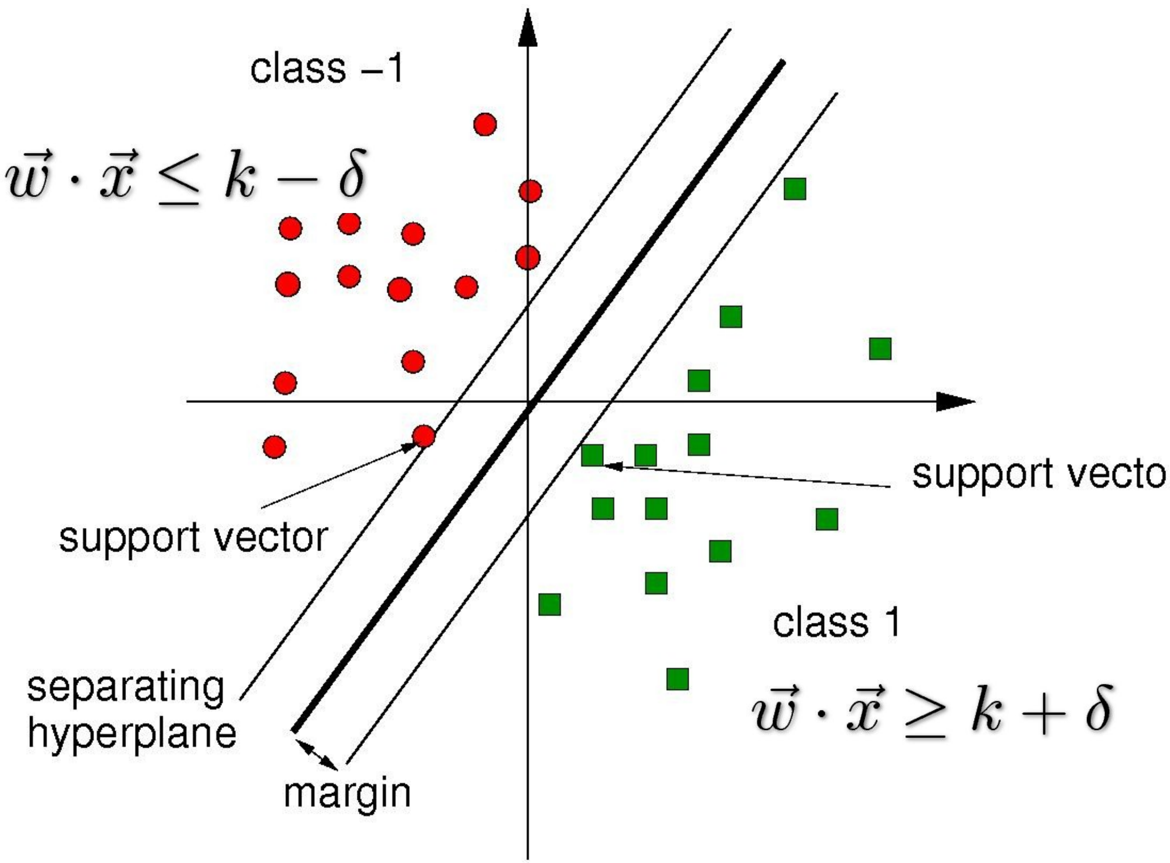

Here the hyperplane is adjusted to maximise the margin

Linear Classification can fail

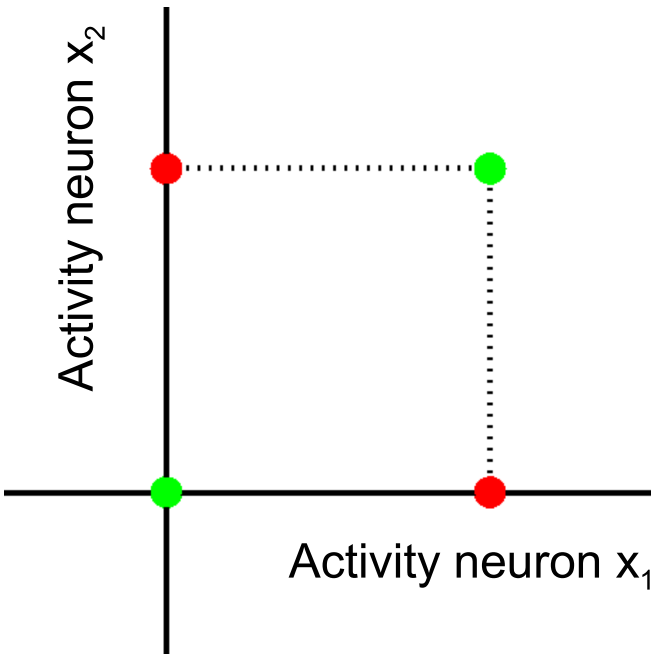

There is no line that can separate green from red.

Dimension = number of neurons

The projection/kernel/hidden layer trick

"Classification is much more probable if the input is first cast into a high-dimensional space by a non-linear transformation."

Cover, T. (1965). "Geometric and statistical properties of systems of linear inequalities with applications in pattern recognition". IEEE T Elect. Comput., 14, 326.

This can be done by using a non-linear "Kernel function" instead of the scalar product $\vec{w} \cdot \vec{x}$.
When used like this it is known as the "**kernel trick**".

M. Aizerman, E. Braverman, and L. Rozonoer (1964). "Theoretical foundations of the potential function method in pattern recognition learning". Automation and Remote Control 25: 821–837

A related concept: MLP

If used with a large hidden layer, multi layer perceptrons (MLP) is a related concept.

See: F. Rosenblatt (1962) “Principles of Neurodynamics”. New York: Spartan books.

Nonlinear projection/ hidden layer/ kernel trick

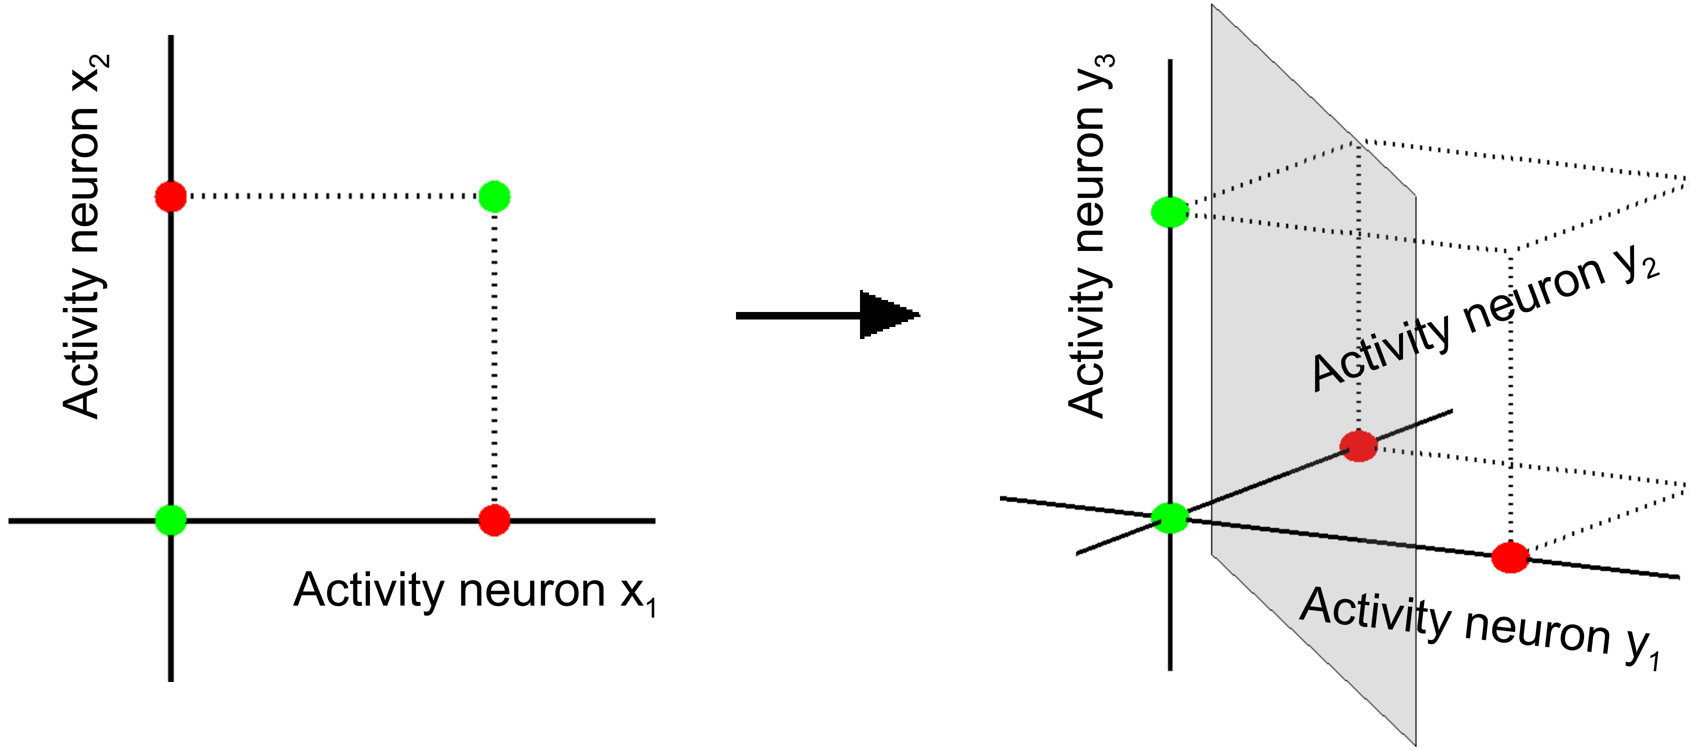

Dimension = number of neurons

Hypothesis: The locust uses this idea

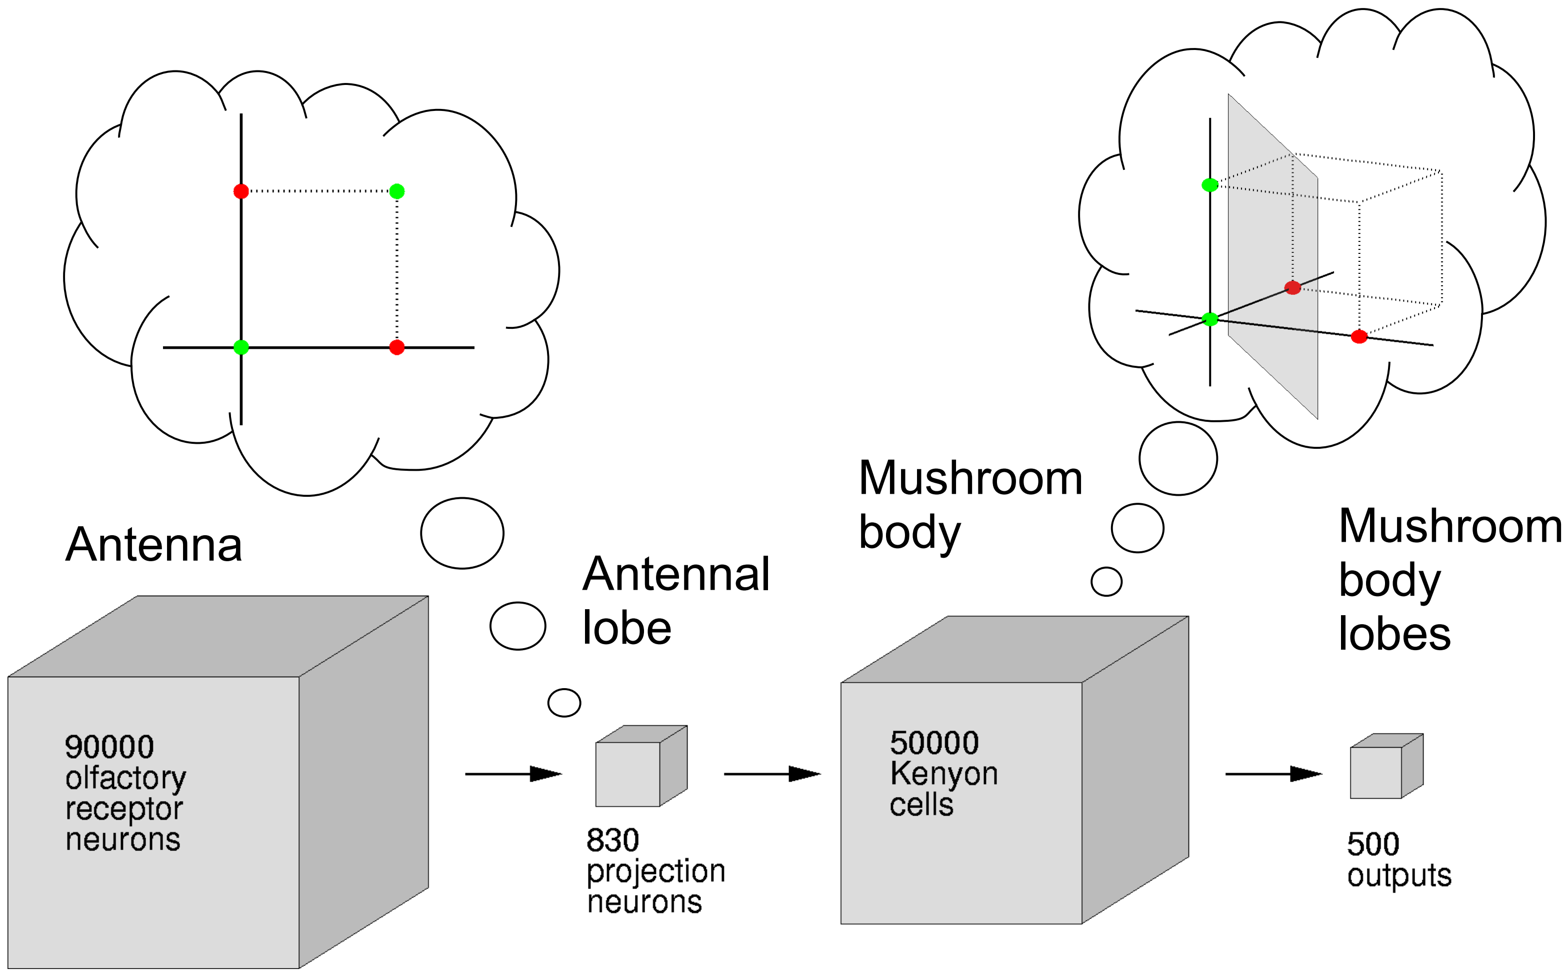

We will use a random connections (a *random kernel,* in a sense)

Dr. Thomas Nowotny,
Centre for Computational Neuroscience and Robotics
University of Sussex

Example result: Classification needs sparse code

… and nature uses it!

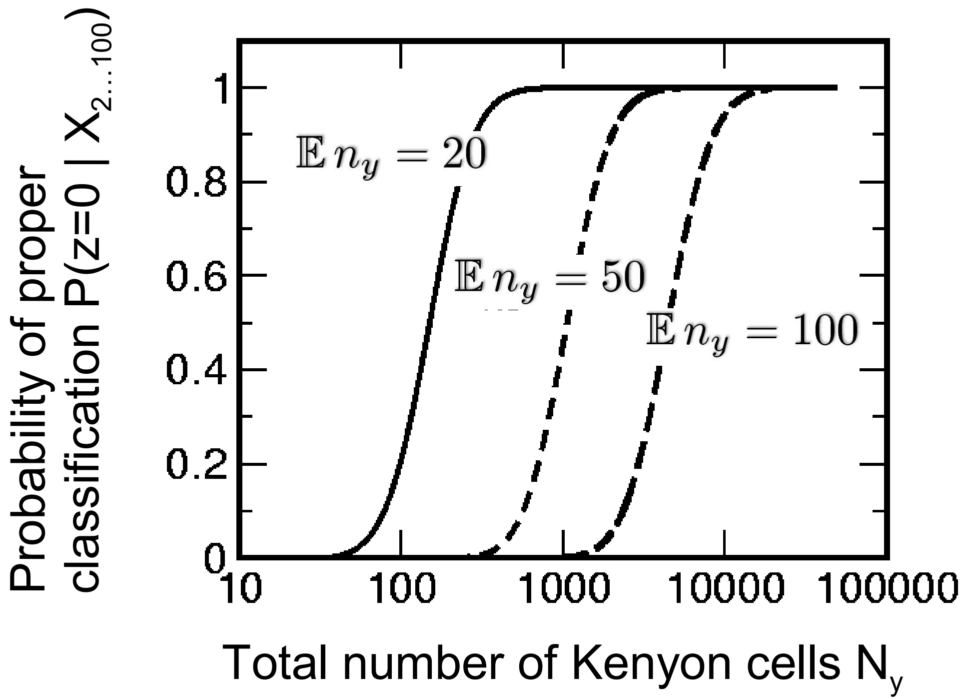

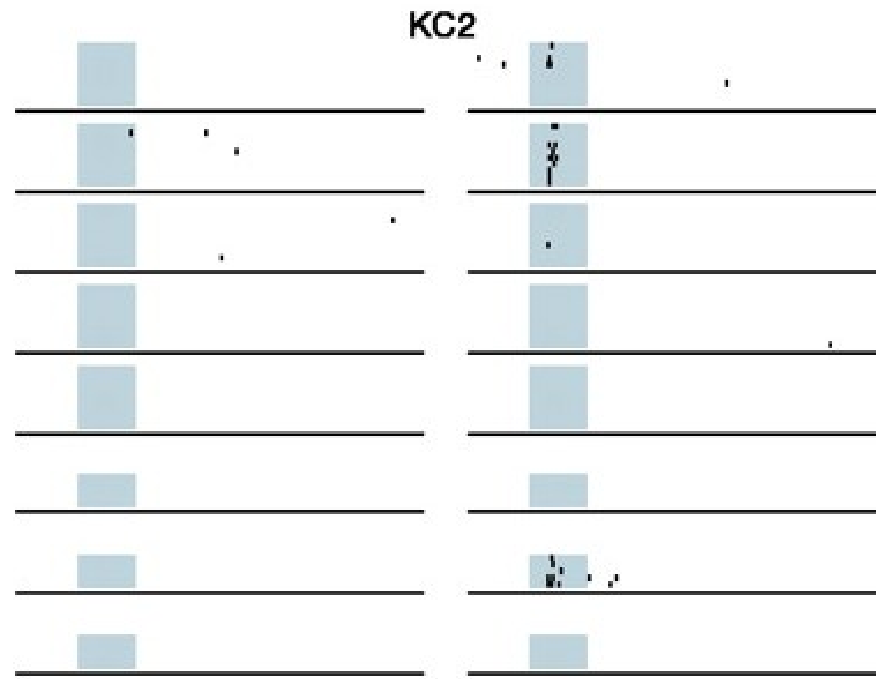

"Have many, but only use a few"

Perez-Orive et al., Science (2002)

Huerta et al., Neural Computation 16(8): 1601-1640 (2004)

Classify multiple classes of inputs

- 10 classes of inputs, 10 patterns each class
- "Winner-take-all" ouputs:
  The output neuron with the strongest input spikes
- Simulations in "*Drosophila* size"

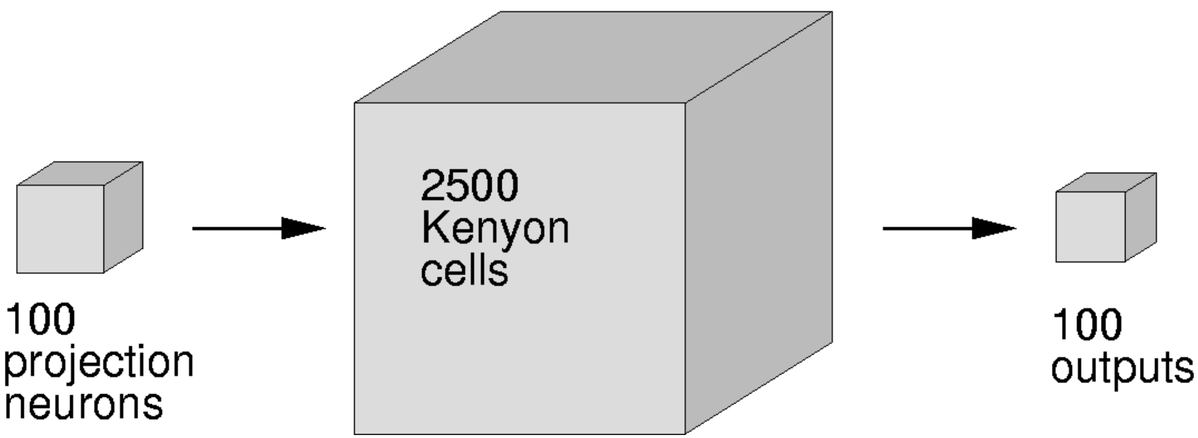

Classes of input patterns

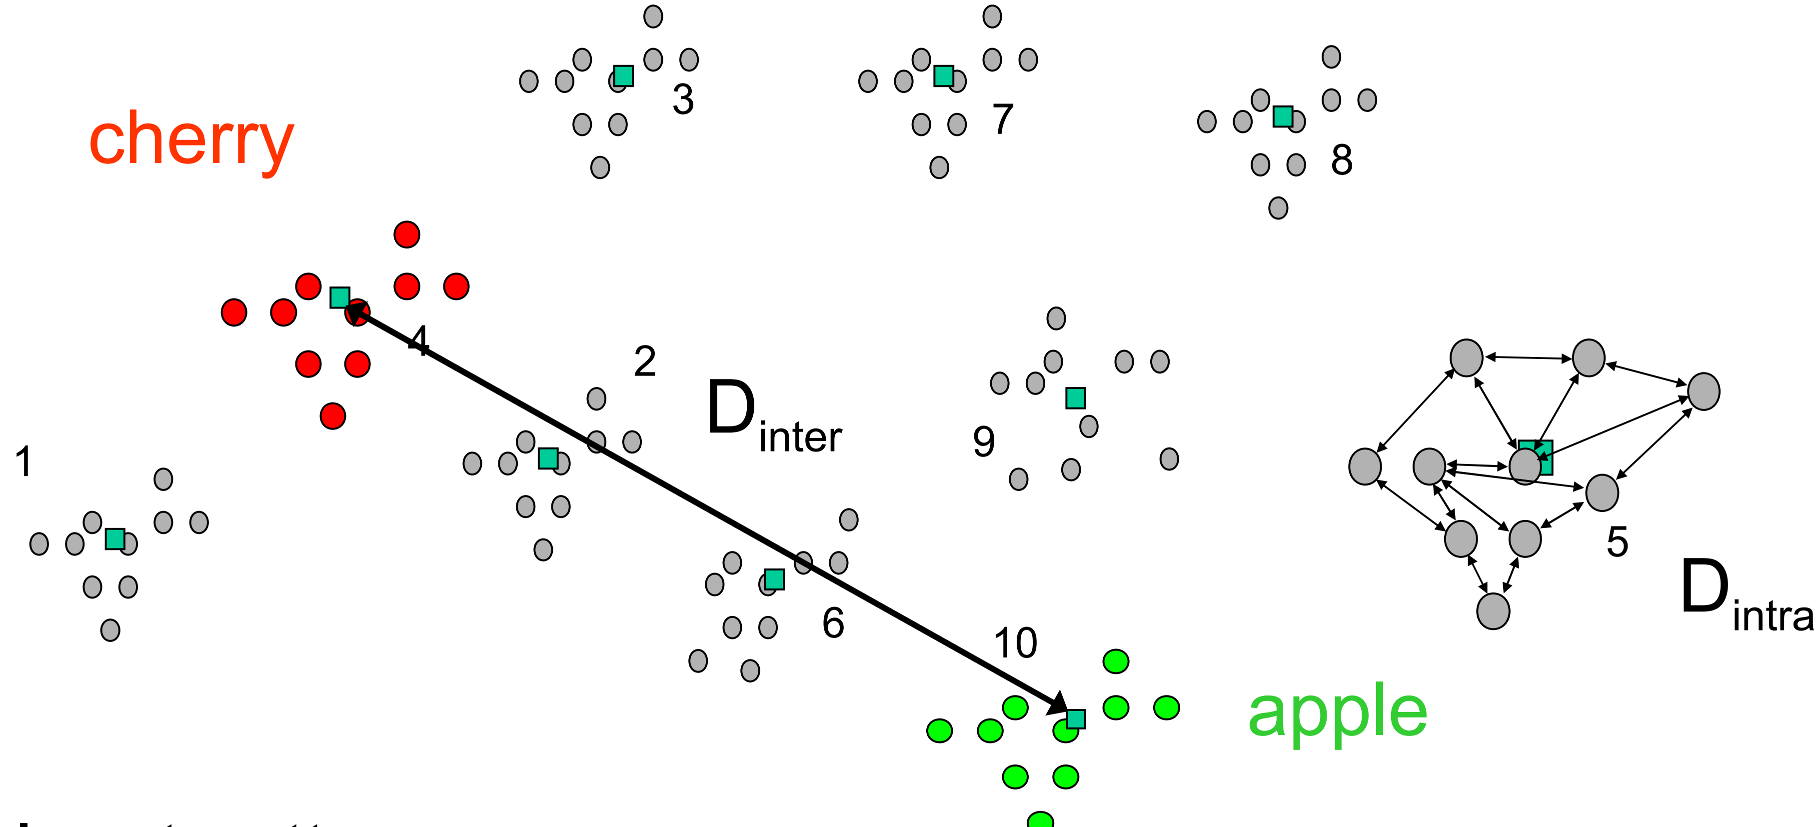

Input pattern space

There are "optimal design parameters"

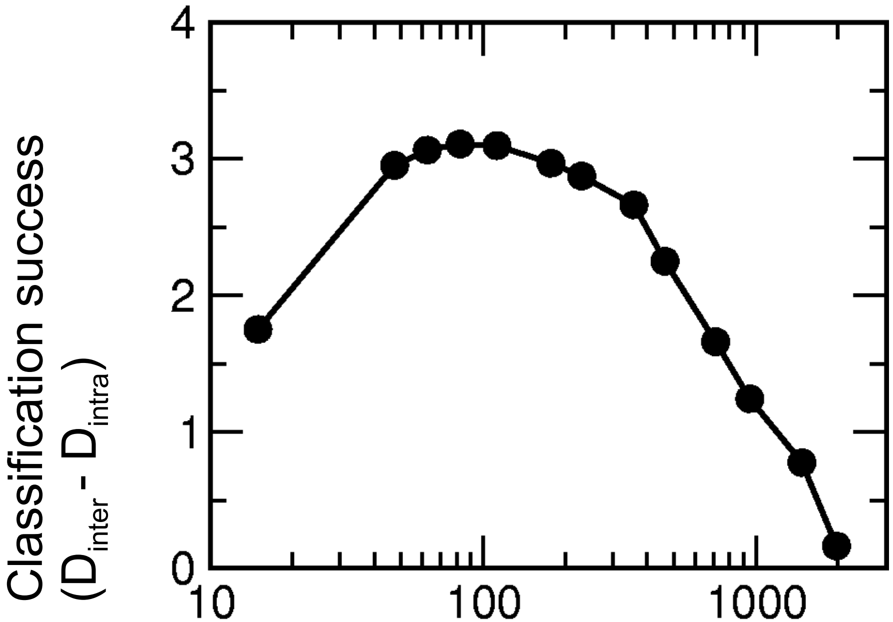

Number of active Kenyon cells $\mathbb{E}\, n_y$

There is an optimal $\mathbb{E}\, n_y$ of active Kenyon cells

Huerta et al., Neural Computation 16 (8):1601-1640 (2004)

Summary

- Random connectivity seems sufficient for classification
- This suggests MLPs with *random kernels* and *local*, “Hebbian” learning
- An optimal, sparse level of activity is postulated *and observed* in biology
- These systems are freely scalable & our analysis provides the parameters of choice
- These systems are extremely robust

Shortcomings

- The winner-take-all competition between output neurons has to be implemented artificially
- Gain control in the MB has to be implemented artificially

These issues can be resolved with more realistic spiking neuron models.

Spiking neuron models

McCulloch-Pitts

Spiking neurons

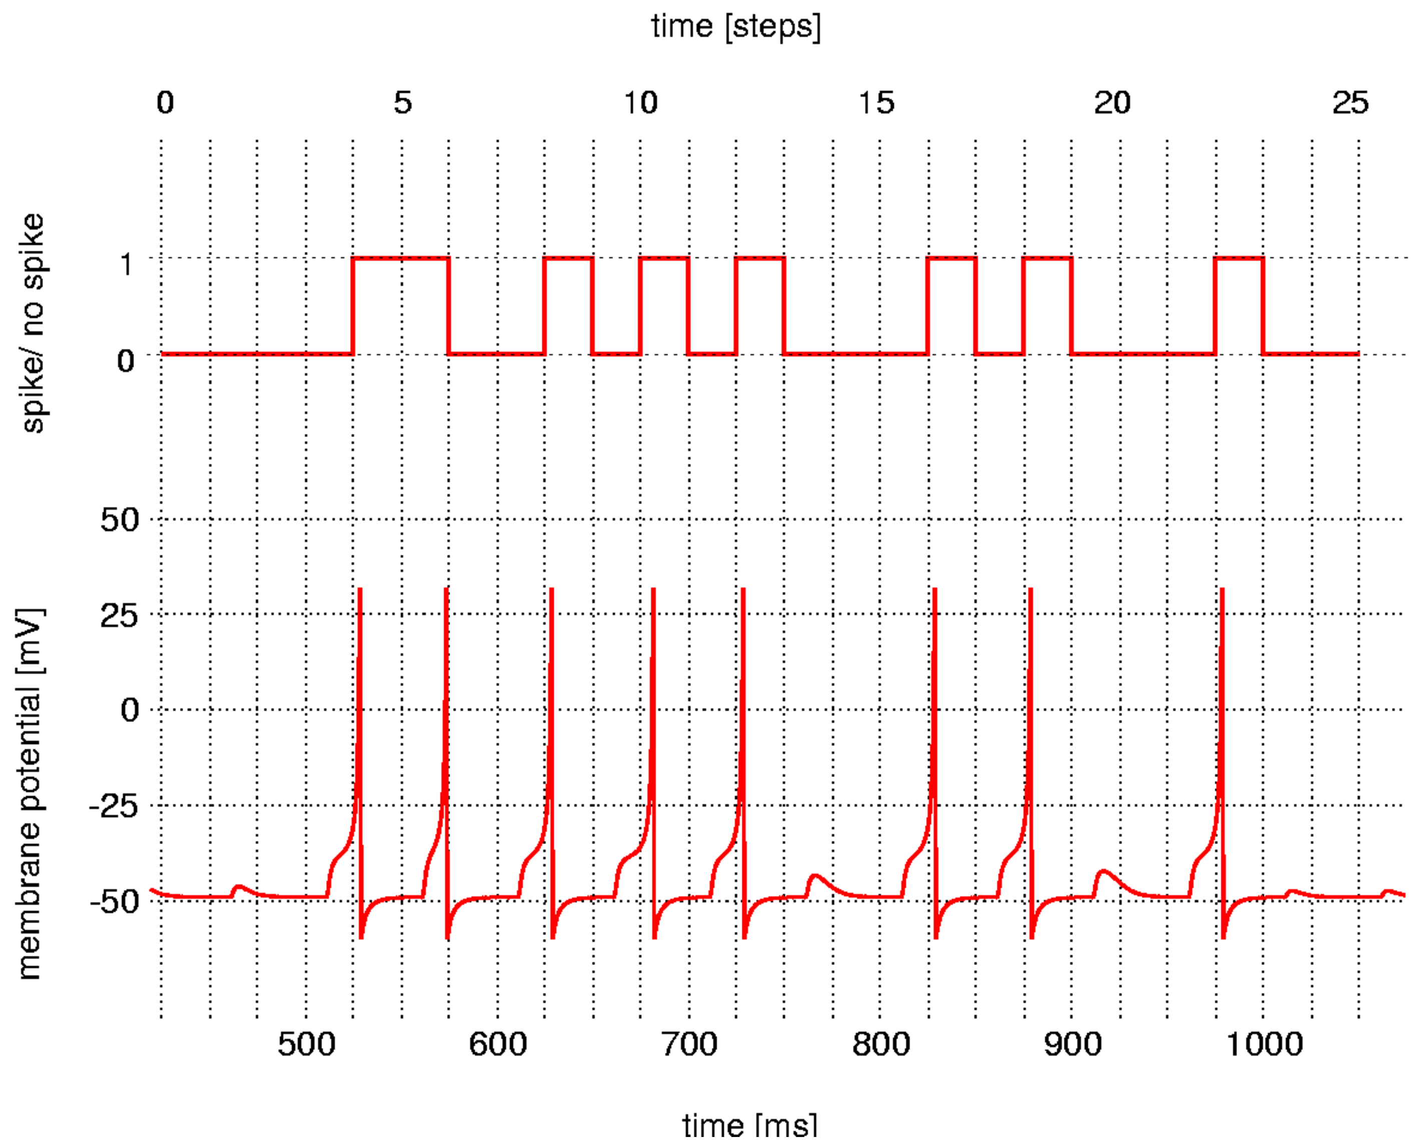

Dr. Thomas Nowotny,
Centre for Computational Neuroscience and Robotics
University of Sussex

Process of recognition: Naïve locust

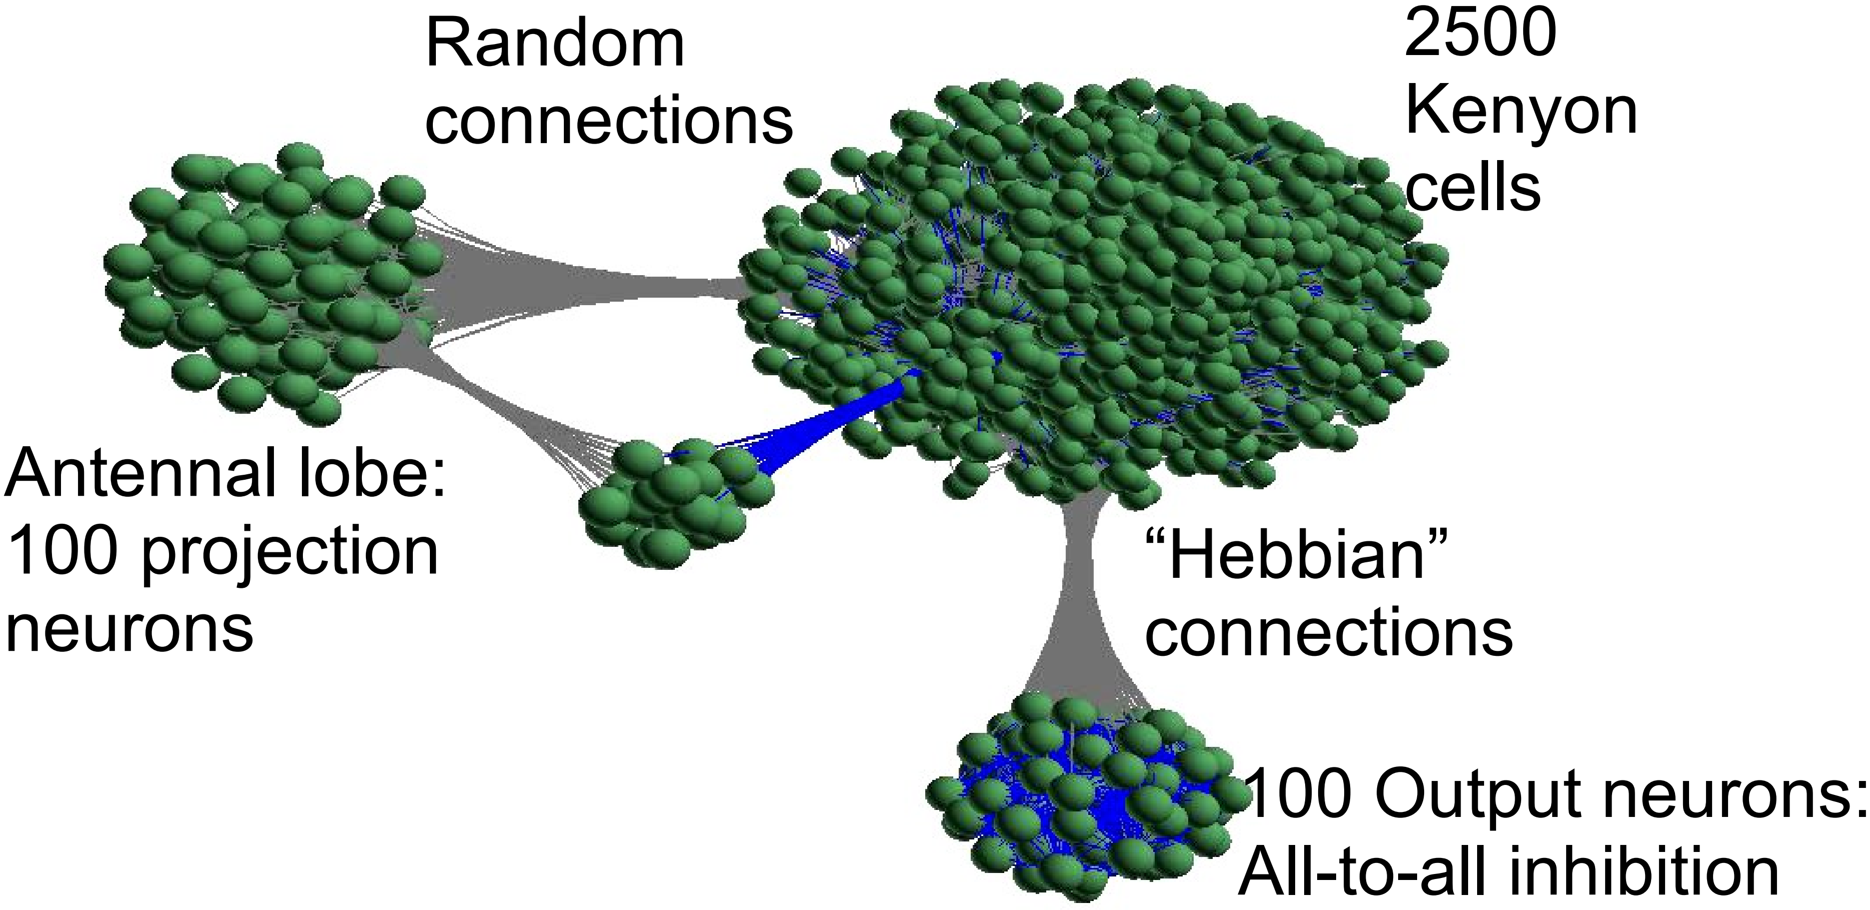

Created with neuranim
http://sourceforge.net/projects/neuranim

Naïve System

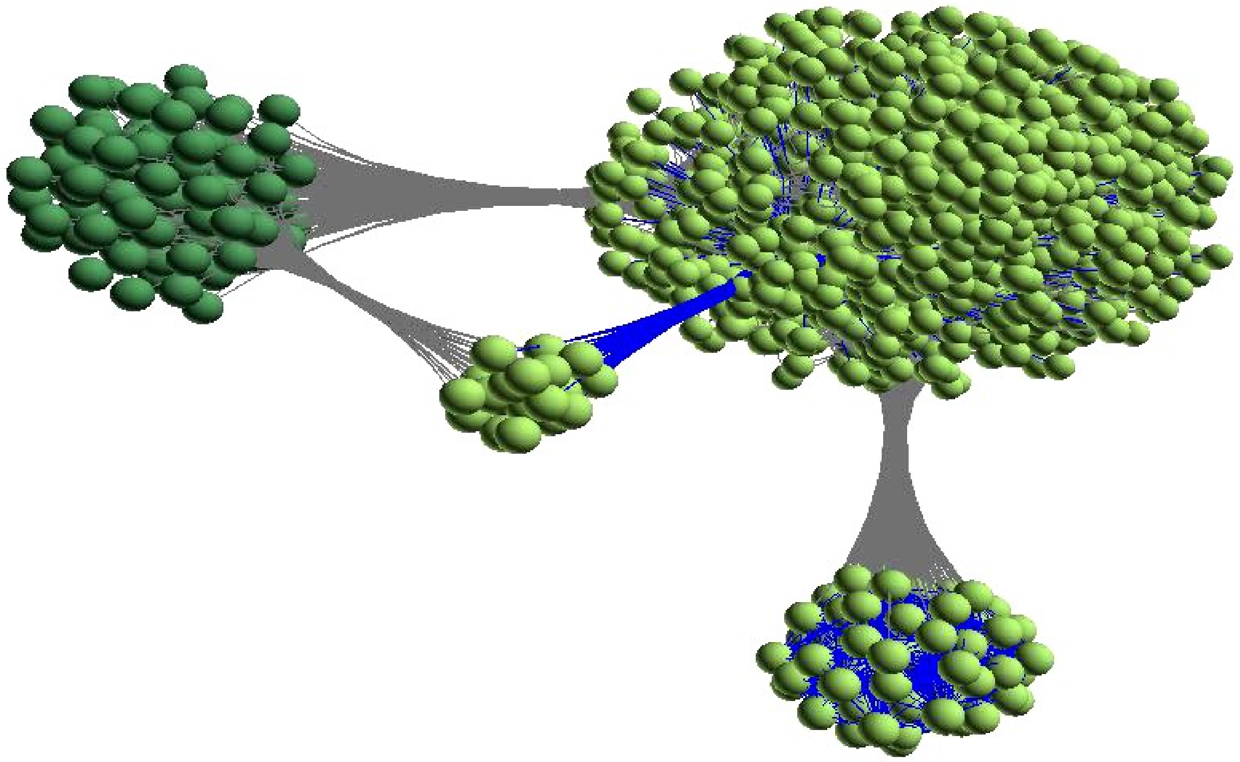

http://sourceforge.net/projects/neuranim

Experienced System

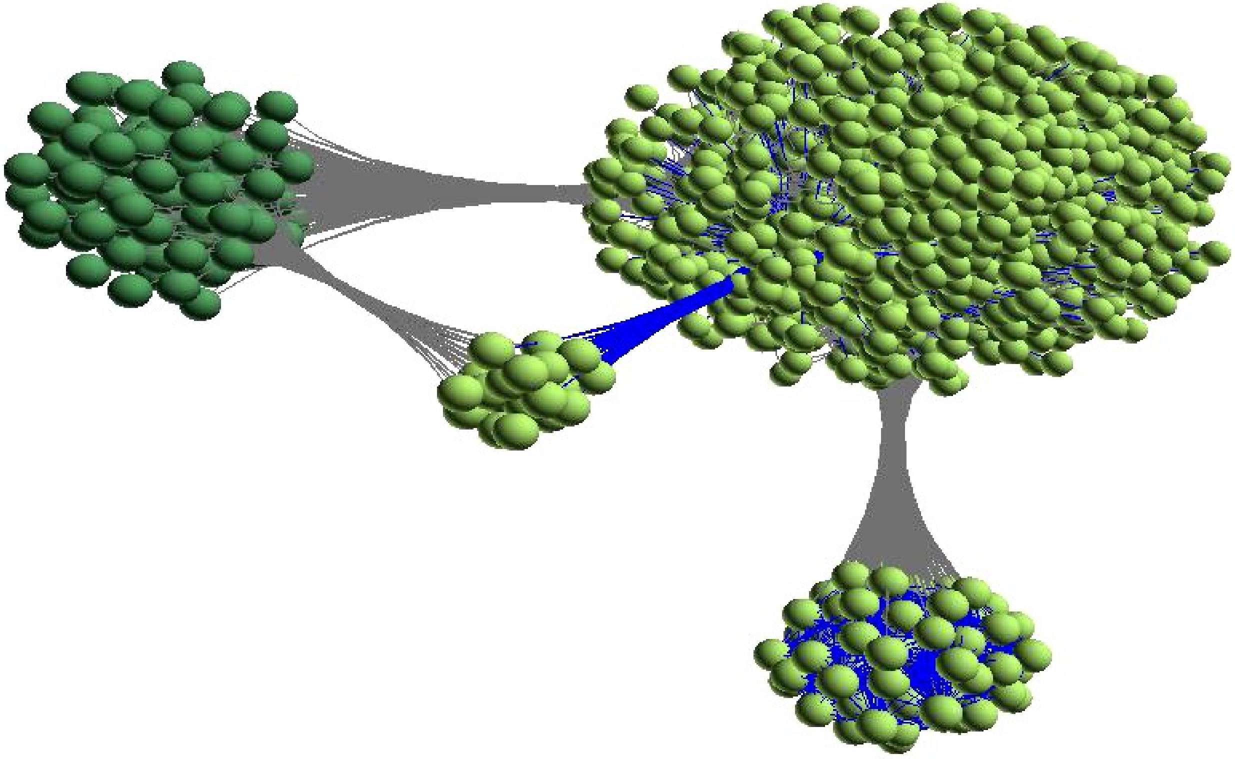

http://sourceforge.net/projects/neuranim

Quantitative Analysis

Antennal lobe

Mushroom body

Naïve system output

Experienced system output

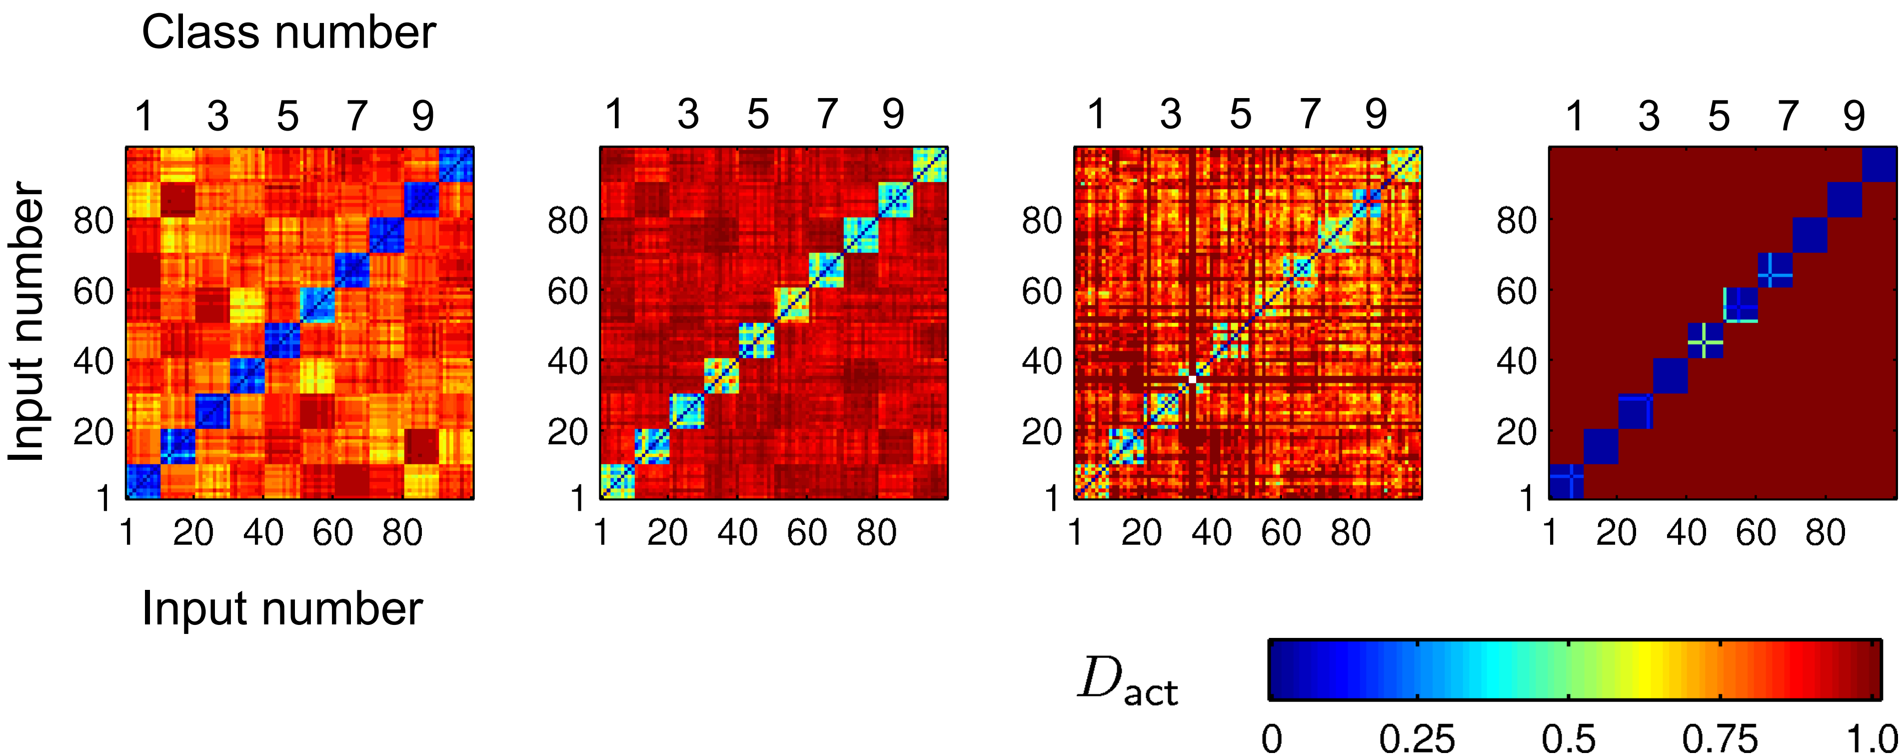

Nowotny et al. Biol Cyber, 93 (6): 436-446 (2005),

Quantitative Analysis

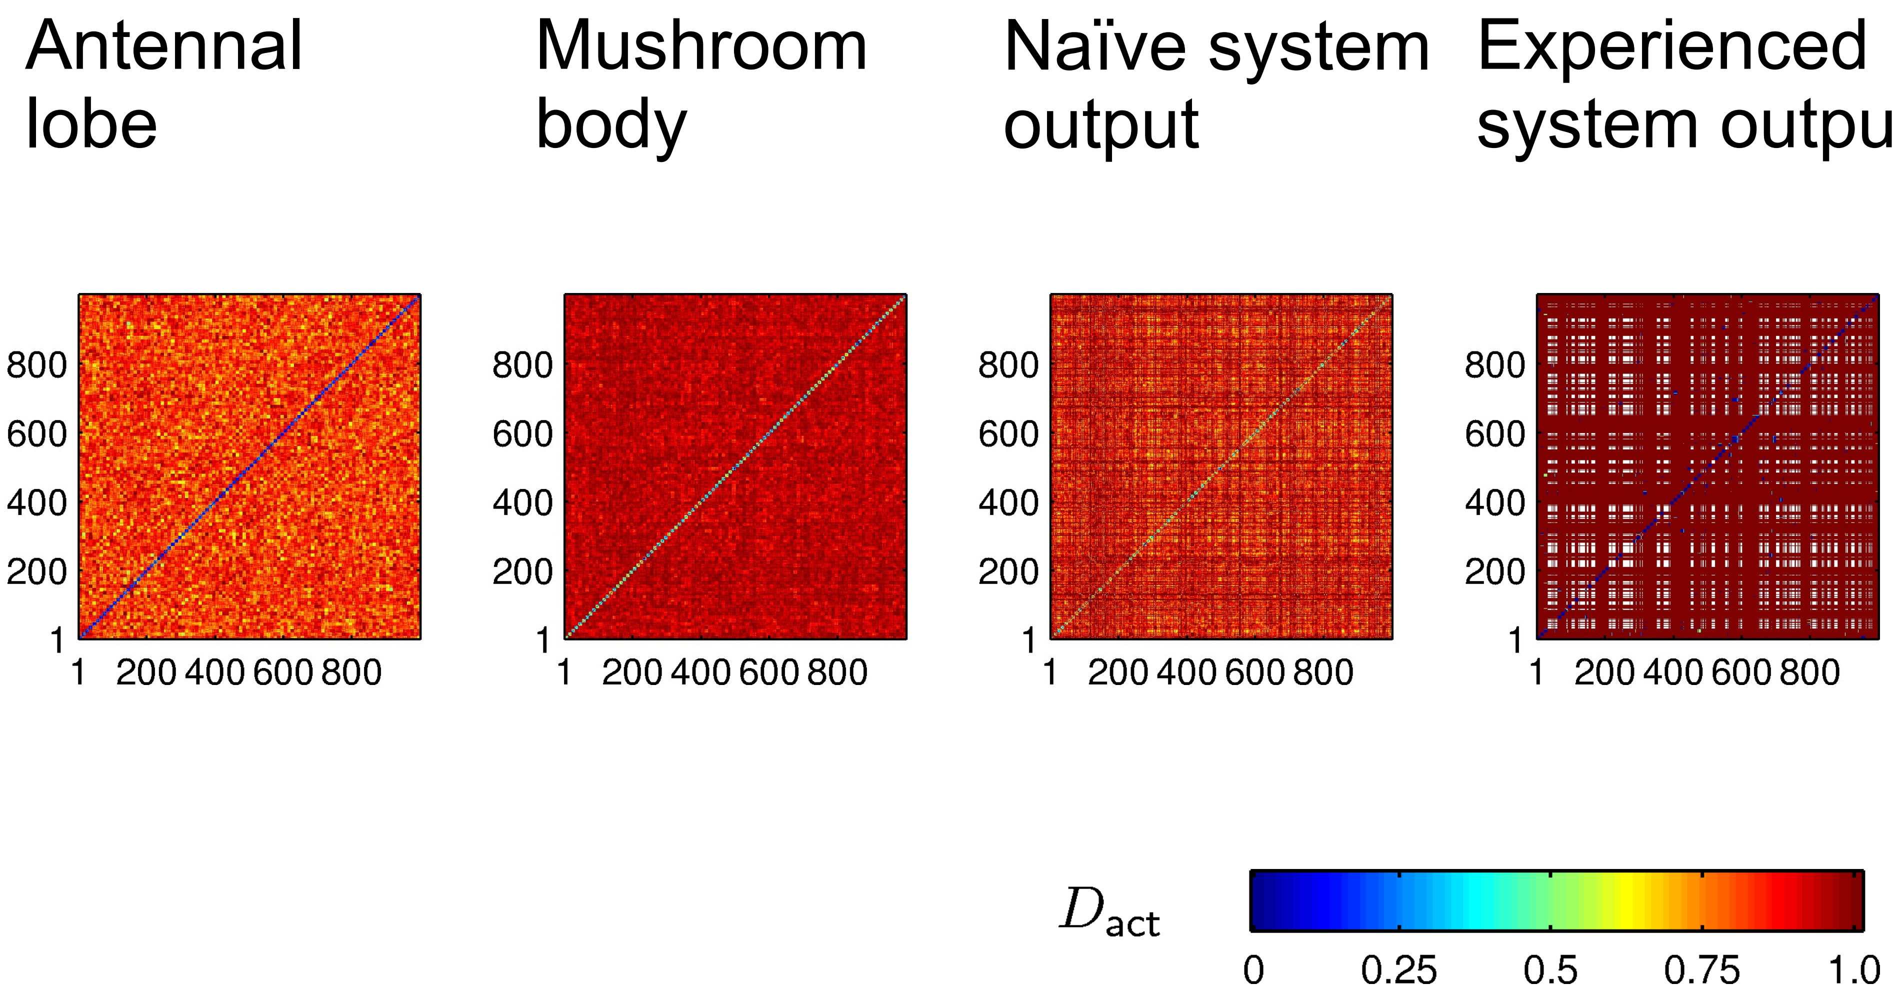

Nowotny et al. Biol Cyber, 93 (6): 436-446 (2005),

Automatic detection of input set structure

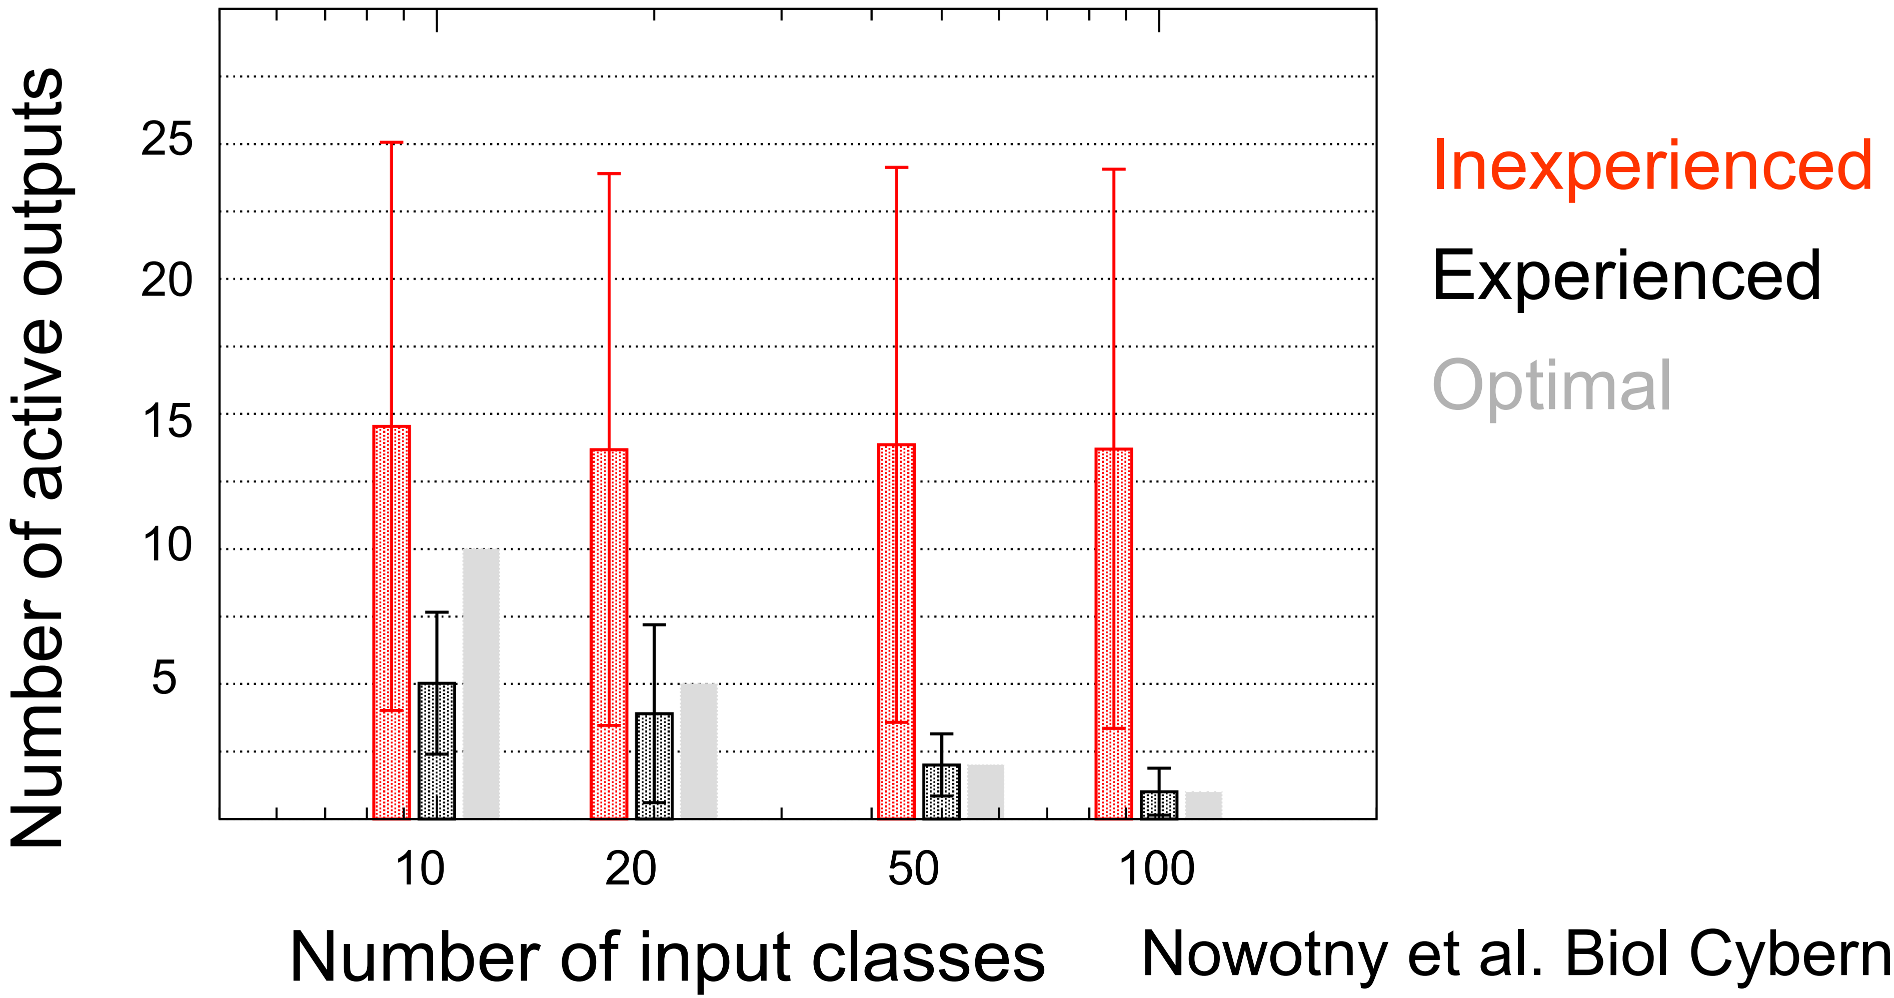

Nowotny et al. Biol Cybern, 93 (6): 436-446 (2005),

Summary: Spiking model

- More realistic biophysical models demonstrate that the system can *self-organize* to recognize odors
- The system detects *the structure of the input pattern set* autonomously

Beyond olfaction ...

- So far I have presented our investigations of the olfactory system of insects using synthetic input data.
- To assess the suggested system as a classifier we applied it to the standard benchmark of the MNIST database of handwritten digits.
- Performance was decent but (as expected) not a revolution in machine learning.

Huerta and Nowotny,
Neural Computation (2009)

Lastly: Implementation of the model(s) on NVidia® CUDA™

Tesla™ S1070 GPU

128 cores, 1.5 GHz

- Split model in to little pieces that execute in parallel
- Take care to use different memory structures in a smart way,

Timing results

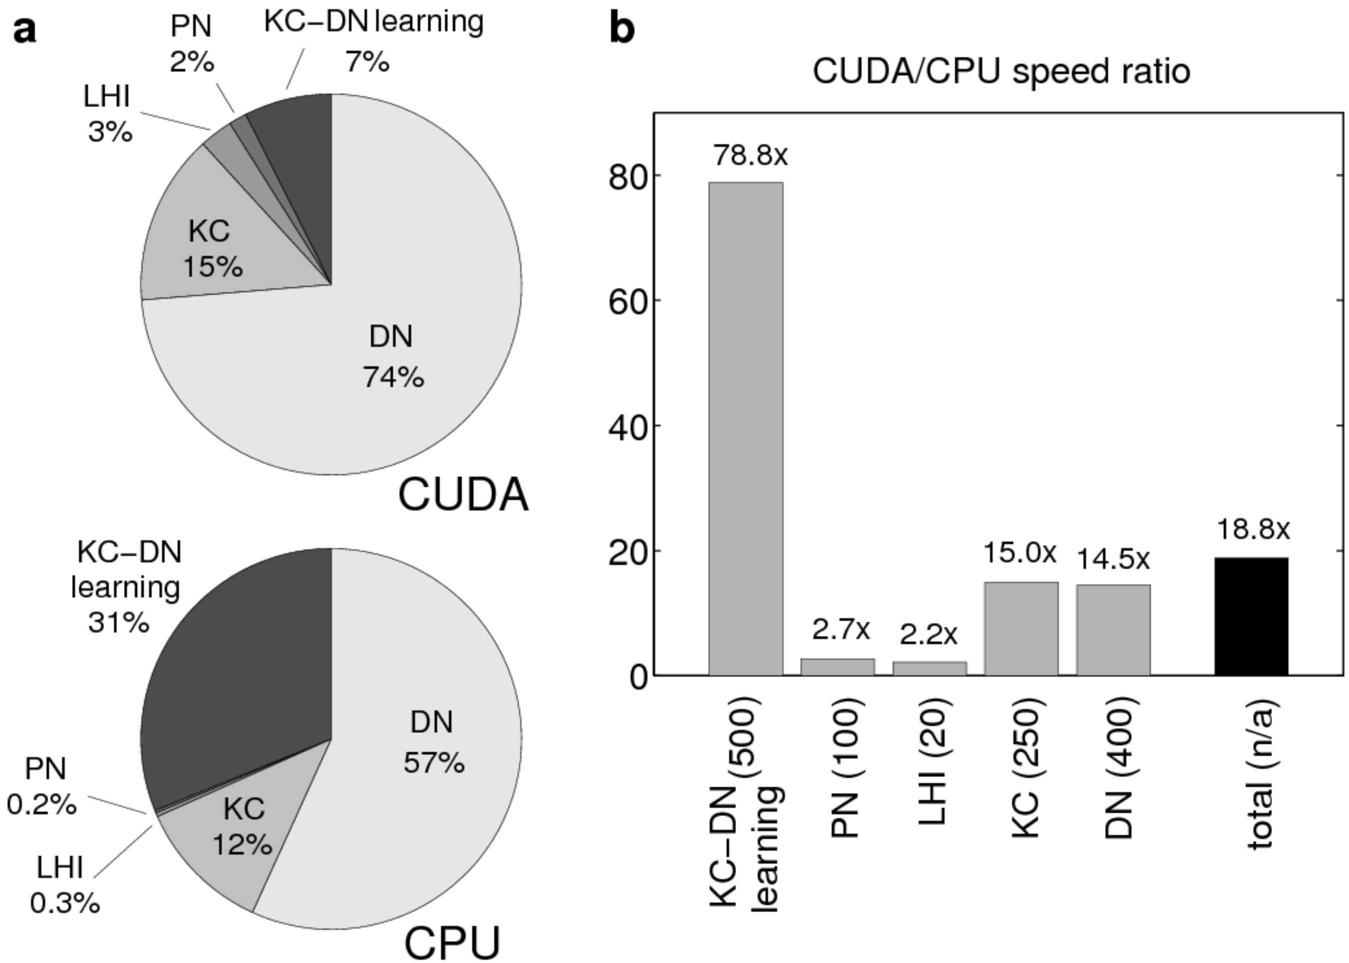

Thomas Nowotny, WCCI Barcelona, 2010

Summary: CUDA™

- I saw an 18 fold speedup on NVidia® Tesla™ S1070 compared to a fast multi-core PC (with AMD® Phenom™ II X4 940 quad core processor at 3 GHz and 4 GB of RAM)
- The systems lends itself to modern prallel architectures due to its feedforward structure.
- However: The details of the implementation do matter (a lot)

Upcoming work

- Pheromone ratio recognition in the moth MGC (with A. Zavada)

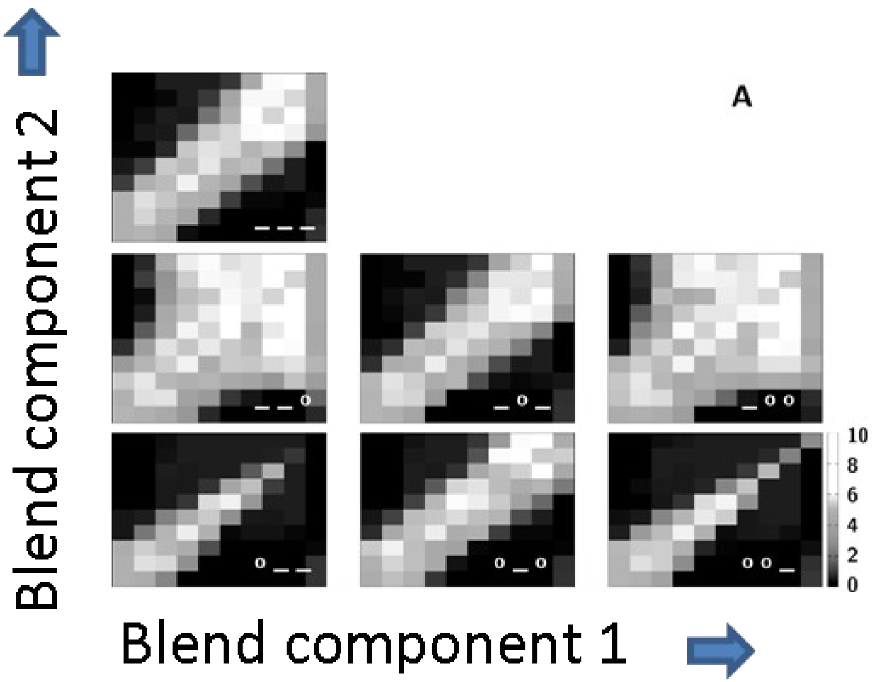

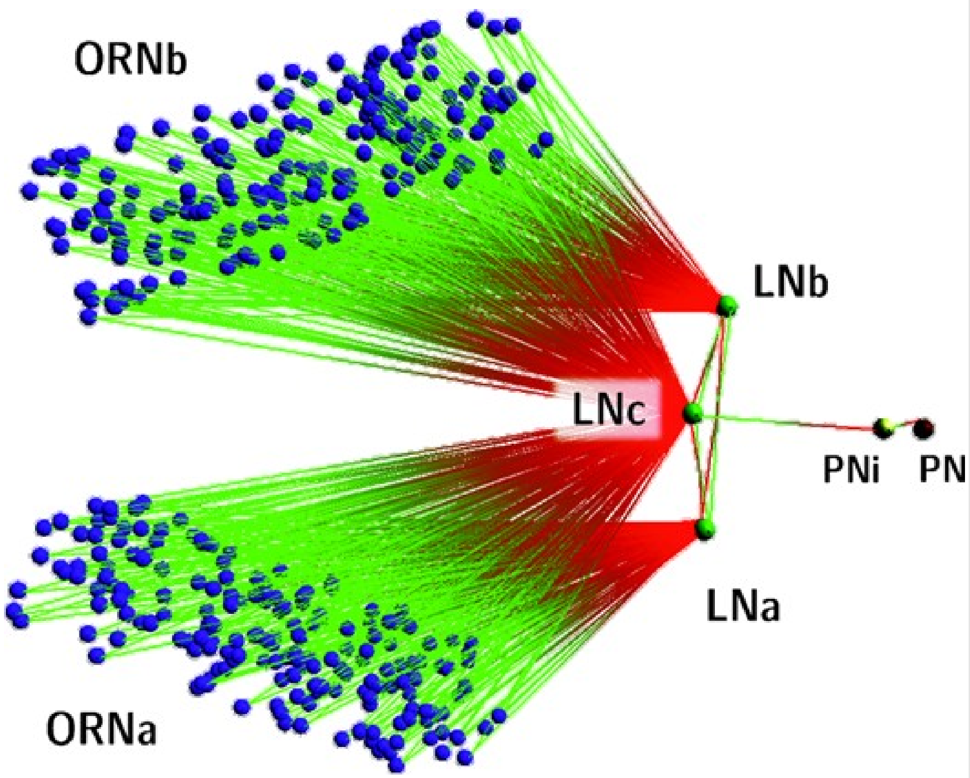

Upcoming work

- Dynamical models of the moth MGC (criticality) (with Chris Buckley)

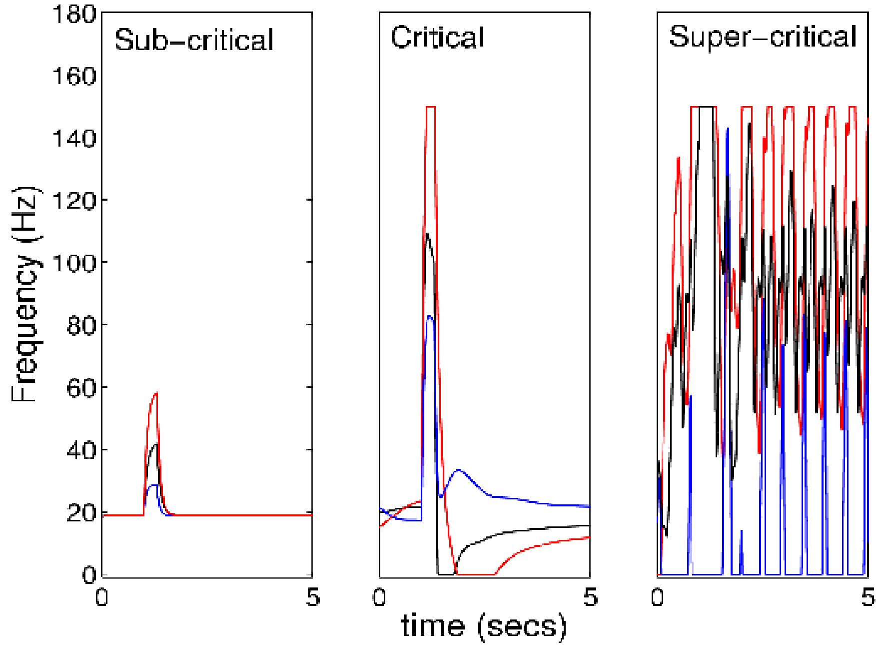

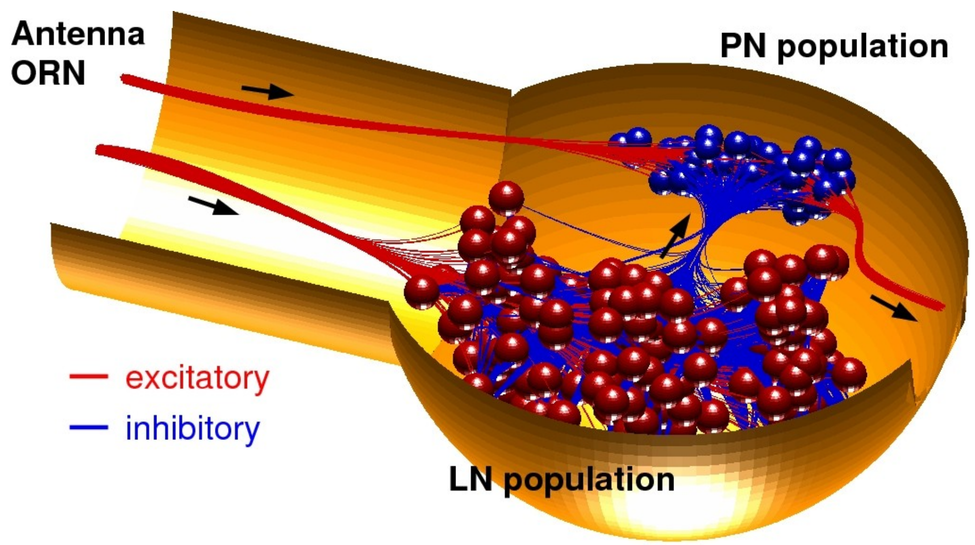

Acknowledgements

R Huerta: Abstract models, early work, MNIST work

M Garzia-Sanchez: Early work on fan-out network

M. Kerem Muezzinoglu: CUDA work

MI Rabinovich: Dynamical aspects

HDI Abarbanel: Initiator of the olfaction program at the INLS

G Laurent: Experimental data, discussions & reality check

A. Zavada & C. L. Buckley: Upcoming MGC work

Thank you for your attention!

Work funded by

Dr. Thomas Nowotny,
Centre for Computational Neuroscience and Robotics

University of Sussex