Deep Neural Networks:

the Quest for Human-Like Brain Power

Iveta Mrázová, Ph.D.

Department of Theoretical Computer Science and Mathematical Logic, Faculty of Mathematics and Physics, Charles University in Prague, Czech Republic

Plenary Talk, CAS 2015, San José, USA, 3. 11. 2015

Outline

- Motivation
- Deep Neural Architectures
- Applications
- Conclusions

Motivation

- Large Government-Initiated Brain Research Projects
- Connectionism Implemented in Silicon
- New Imaging Technologies
- New Neurobiological Discoveries

Large-Scale Brain Projects

the WHITE HOUSE
BRAIN INITIATIVE
BRAIN RESEARCH THROUGH ADVANCING INNOVATIVE NEUROTECHNOLOGIES

BRAIN Initiative - USA, April 2013

~ **B**rain **R**esearch Through **A**dvancing **I**nnovative **N**eurotechnologies

on a par with the Apollo Program to land humans on the moon
http://www.whitehouse.gov/infographics/brain-initiative

Goal: understand the human mind and uncover new ways to treat, prevent, and cure brain disorders like Alzheimer’s, schizophrenia, autism, epilepsy, and traumatic brain injury

Expected costs: > 4 billion USD / 10 years

Participants:
- **DARPA** ~ Defense Advanced Research Projects Agency
- **IARPA** ~ Intelligence Advanced Research Projects Activity
- **NIH** ~ National Institutes of Health
- **NSF** ~ National Science Foudation
- **FDA** ~ Food and Drug Administration
- **private sector**

Large-Scale Brain Projects

Human Brain Project

HBP ~ **H**uman **B**rain **P**roject – EU, January 2013
https://www.humanbrainproject.eu/

Goal: understand what makes us human (through brain–wide analyses of neural network activity at the level of single neurons), develop new treatments for brain disorders and build revolutionary new computing technologies.

13 Subprojects: *Strategic Mouse Brain Data*
Strategic Human Brain Data
The Brain Simulation Platform,
The High Performance Computing Platform, etc.

Expected costs: 1.3 billion USD /10 year
Participants: **112 partners from 24 countries**

Large-Scale Brain Projects

Brain/MINDS Project - Japan, 2014

~ Brain Mapping by Integrated Neurotechnologies for Disease Studies
http://brainminds.jp/en/

Goal: study the neural networks controlling higher brain functions in the marmoset, to get new insights into information processing and diseases of the human brain such as dementia and depression

Expected costs: 300 million USD / 10 years

Participants: **RIKEN Brain Science Institute** – Core Institute
Keio University – Partner Institute
Kyoto University – Partner Institute
… and several other institutions (mainly academic)

Large-Scale Brain Projects

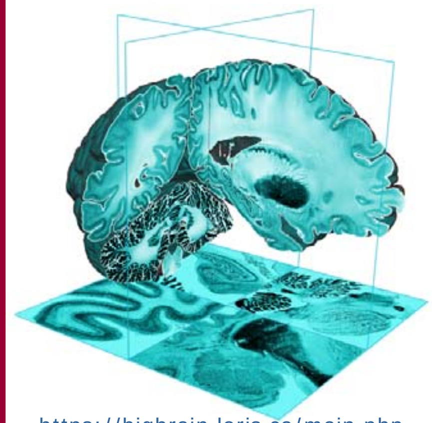

Government-Initiated:

- **China Brain Science Project**, 2015
  - is focused on developmental, psychiatric and neurodegenerative disorders and should promote breakthroughs in AI research to reshape country´s industry, military, and service structure for the new industrial revolution
- huge projects launched also by **Israel** and **Canada**

Other brain research projects include:

- **Allen Brain Atlas** - Allen Institute for Brain Science, USA, 2003
- **BigBrain** - Montreal Neurological Institute and German Forschungszentrum Jülich, June 2013

  https://bigbrain.loris.ca/main.php

Connectionism Implemented in Silicon: Early Attempts

"ACE" MAY BE FASTEST BRAIN

BRITISH ROBOT ON DISPLAY

DAILY TELEGRAPH REPORTER

An electronic "brain," which is expected to outshine all rivals by its speed in working out mathematical problems, is being developed by the National Physical Laboratory. It is known as "Ace" (automatic computing engine).

One of Ace's 43 "brain cells," 6ft high, was displayed in the library of the Royal Society, Burlington House, yesterday. It was an exhibit in a collection illustrating the development of the National Physical Laboratory, which celebrates its jubilee this year.

Dr. E. C. BULLARD, director of the laboratory told me he hoped that Ace would be completed, with "memory" built in, by the summer. It would then tackle calculations a thousand times as quickly as a girl with a desk computor, and would be able to "remember" 256 10-digit numbers at a time.

Ace should surpass the world's most advanced electronic calculator, completed at Cambridge University mathematical laboratory last summer. It should prove invaluable to scientists engaged on research into atomic energy or aero-dynamics.

Young demonstrators operated yesterday a test panel as easily as if it had been a cricket score-board. But they admitted that Ace could not test Prof. Einstein's latest formulæ. "Ace does not deal with theories—only with practicalities."

The Daily Telegraph, 31 January 1950

Connectionism Implemented in Silicon: The Mark I Perceptron

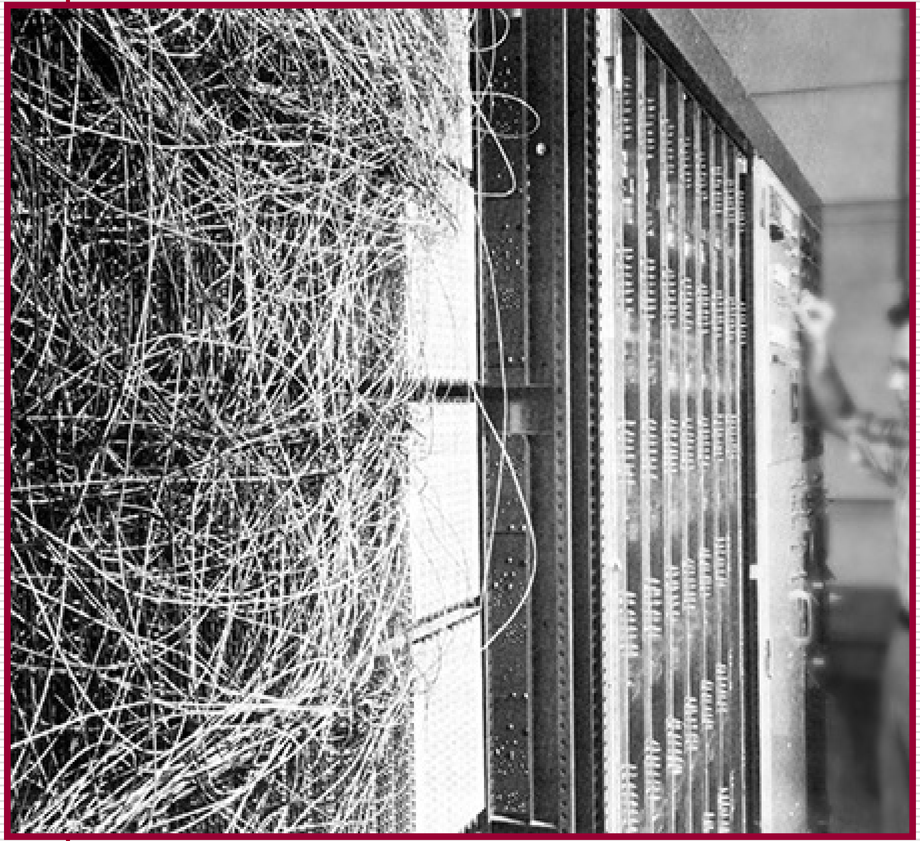

A visual pattern classifier:

- **20x20** photosensitive **input units** modeling a small retina
- **512 hidden units** (stepping motors) each of which could take several excitatory and inhibitory inputs
- **8 output** (response) **units**
- connections from the input to the hidden layer could be altered through **plug-board wiring**, but once wired they remained fixed for the experiment
- connections from the hidden to the output layer were adjusted through **perceptron training**

The Mark I Perceptron, Cornell Aeronautical Laboratory, 1957-1959

Connectionism Implemented in Silicon – the project SyNAPSE

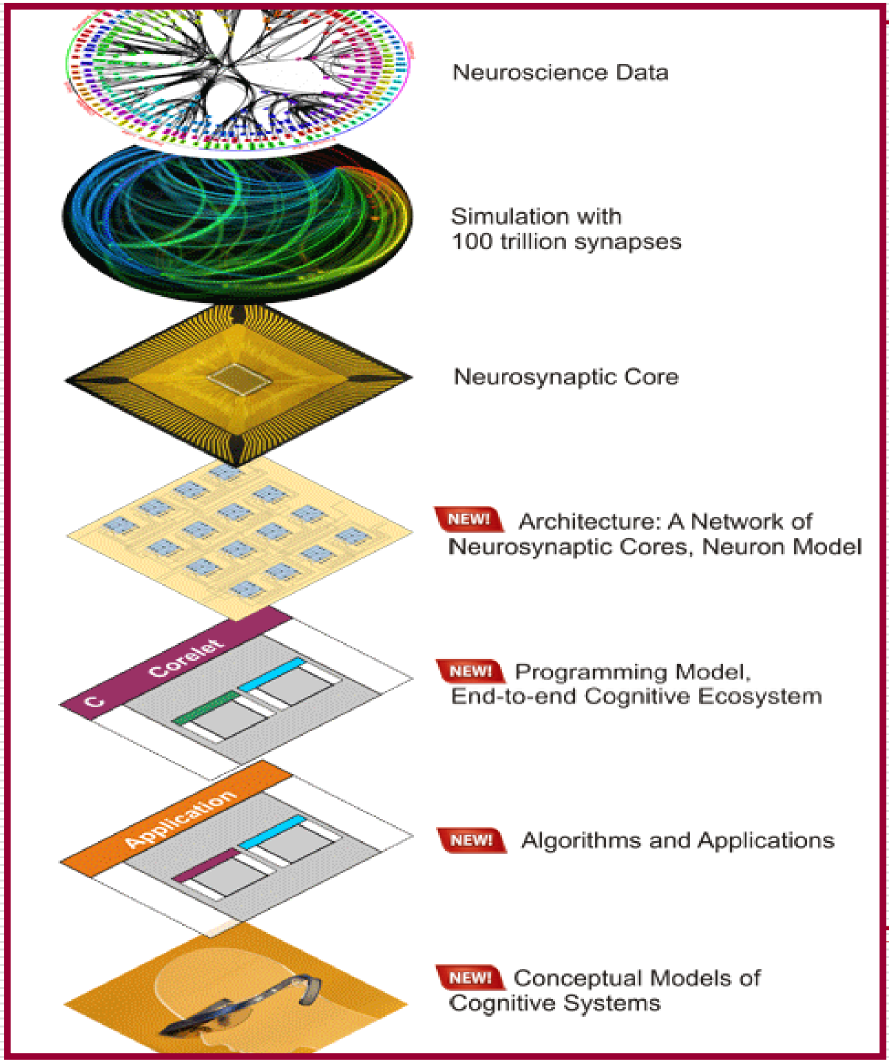

~ **Systems of Neuromorphic Adaptive Plastic Scalable Electronics**

- A DARPA program undertaken by HRL, HP and IBM (Dr. D. Modha)
- **Goal:** develop a novel cognitive computing architecture inspired by the function, low power, and compact volume of the brain
- non von Neumann architecture (neuromorphic computing)
- applications, e.g., in image and video processing, NLP, composer recognition, collision avoidance

Connectionism Implemented in Silicon – Neurosynaptic Chips

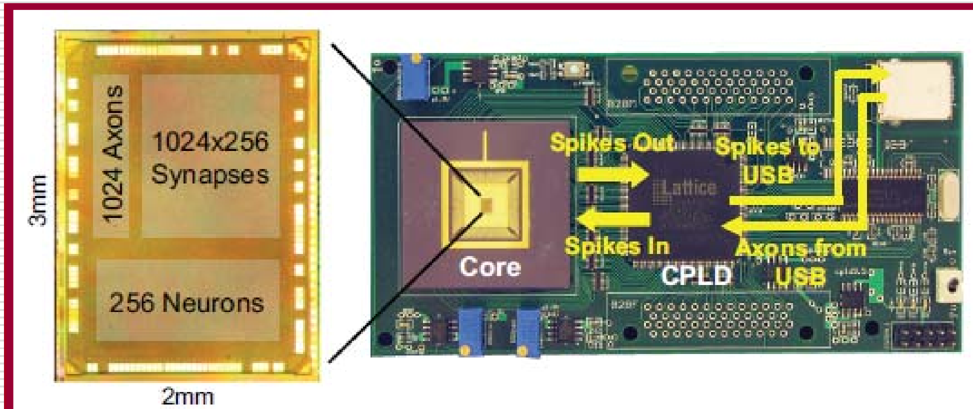

a neurosynaptic core

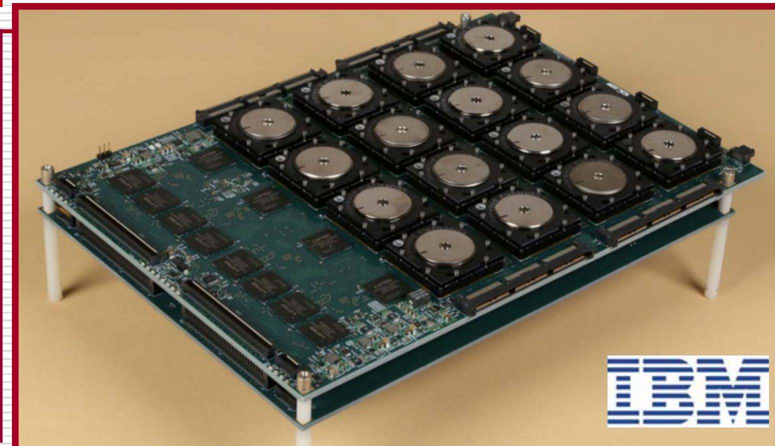

a circuit board with a 4×4 array of SyNAPSE-developed chips

- neurosynaptic TrueNorth Chip (with 4096 neurosynaptic cores)
- 1 million programmable neurons (cca 86 bn in human brains)
- 256 million configurable synapses (cca 10^{14}–10^{15} for humans)
- efficient, scalable, flexible

Connectionism Implemented in Silicon – Data Storage

A new hand-sized tape cartridge can store 220 TB of data:

- big data
- cloud computing
- cheap

- IBM, Sony, …

IBM's Tale of the Tape

More than 60 years of tape innovation

	2006	2010	2014	2015
Aerial Density (bits per sq inch)	6.67 Billion	29.5 Billion	85.9 Billion	123 Billion
Cartridge Capacity	8 Terabytes	35 Terabytes	154 Terabytes	220 Terabytes
Number of Books Stored	8 Million	35 Million	154 Million	220 Million
Track Width (micrometers)	1.5	0.45	0.177	0.140
Linear Density (bits per inch)	400'000	518'000	600'000	680'000
Tape Material	Barium Ferrite	Barium Ferrite	Barium Ferrite	Barium Ferrite
Tape Thickness (micrometers)	6.1	5.9	4.3	4.3
Tape Length (meters)	890	917	1255	1255

New Imaging Technologies:

Lightsheet Microscopy

Z1 was world´s first program-controlled computer

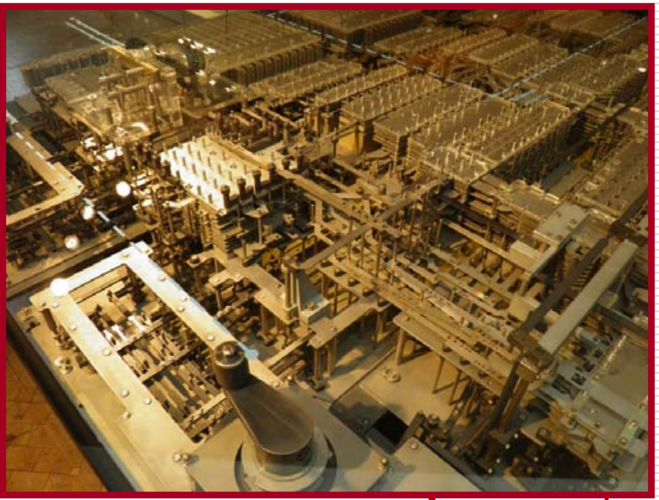

- **ZEISS Lightsheet Z.1:**
  - 32 TB Storage and Data Analysis Module
  - weights cca 500 lbs

New Imaging Technologies:

Lightsheet Microscopy

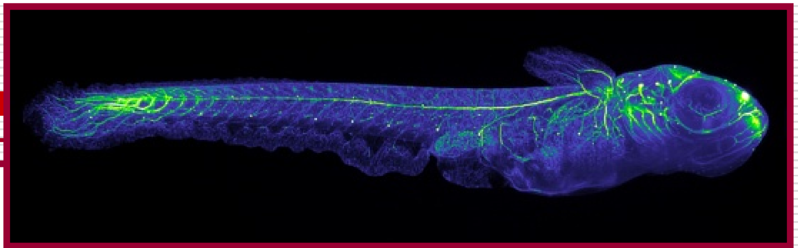

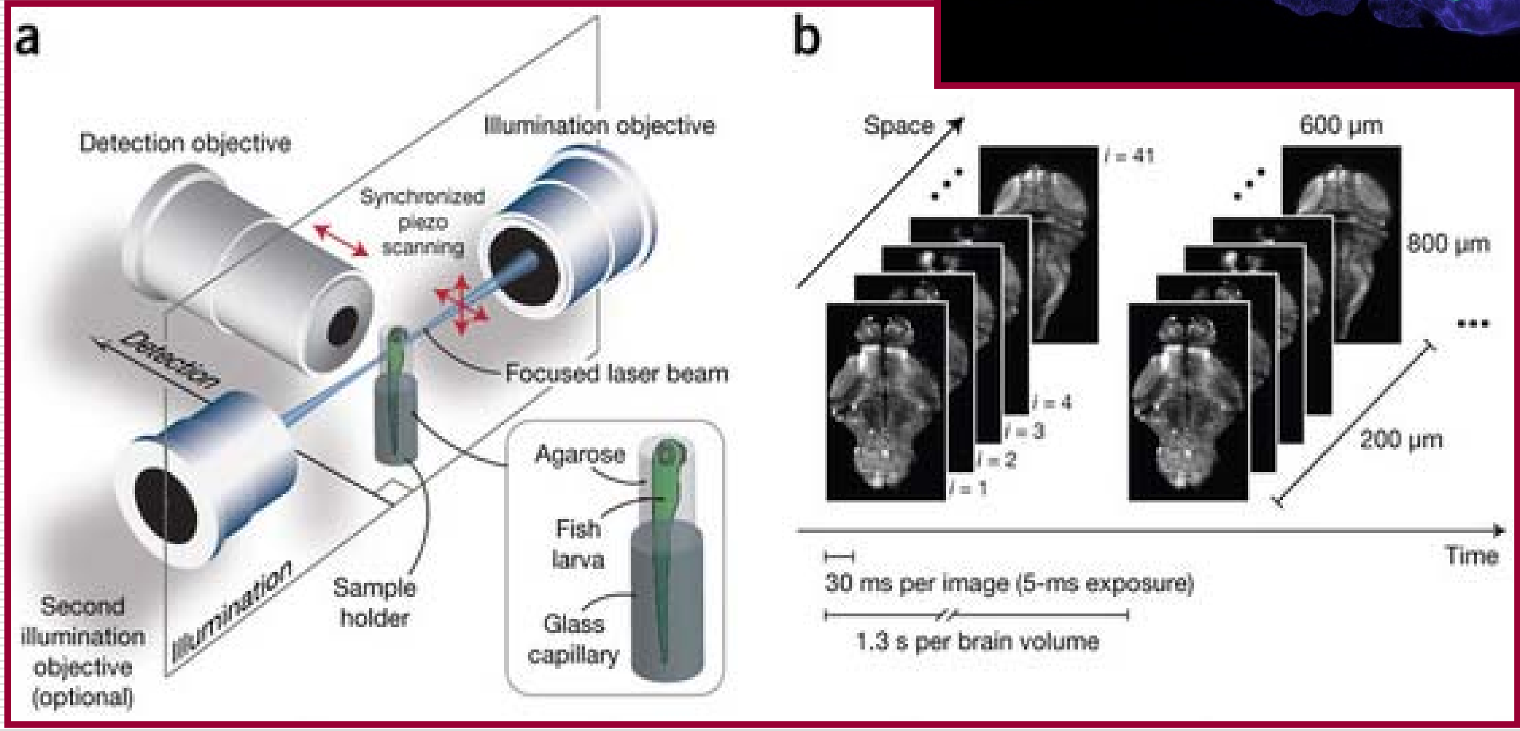

- based on the principles of ultramicroscope developed by Richard Adolf Zsigmondy in 1902 (Nobel Prize in 1925)

http://www.nature.com/nmeth/journal/v10/n5/fig_tab/nmeth.2434_SV4.html

New Imaging Technologies:
Lightsheet Microscopy – System Requirements

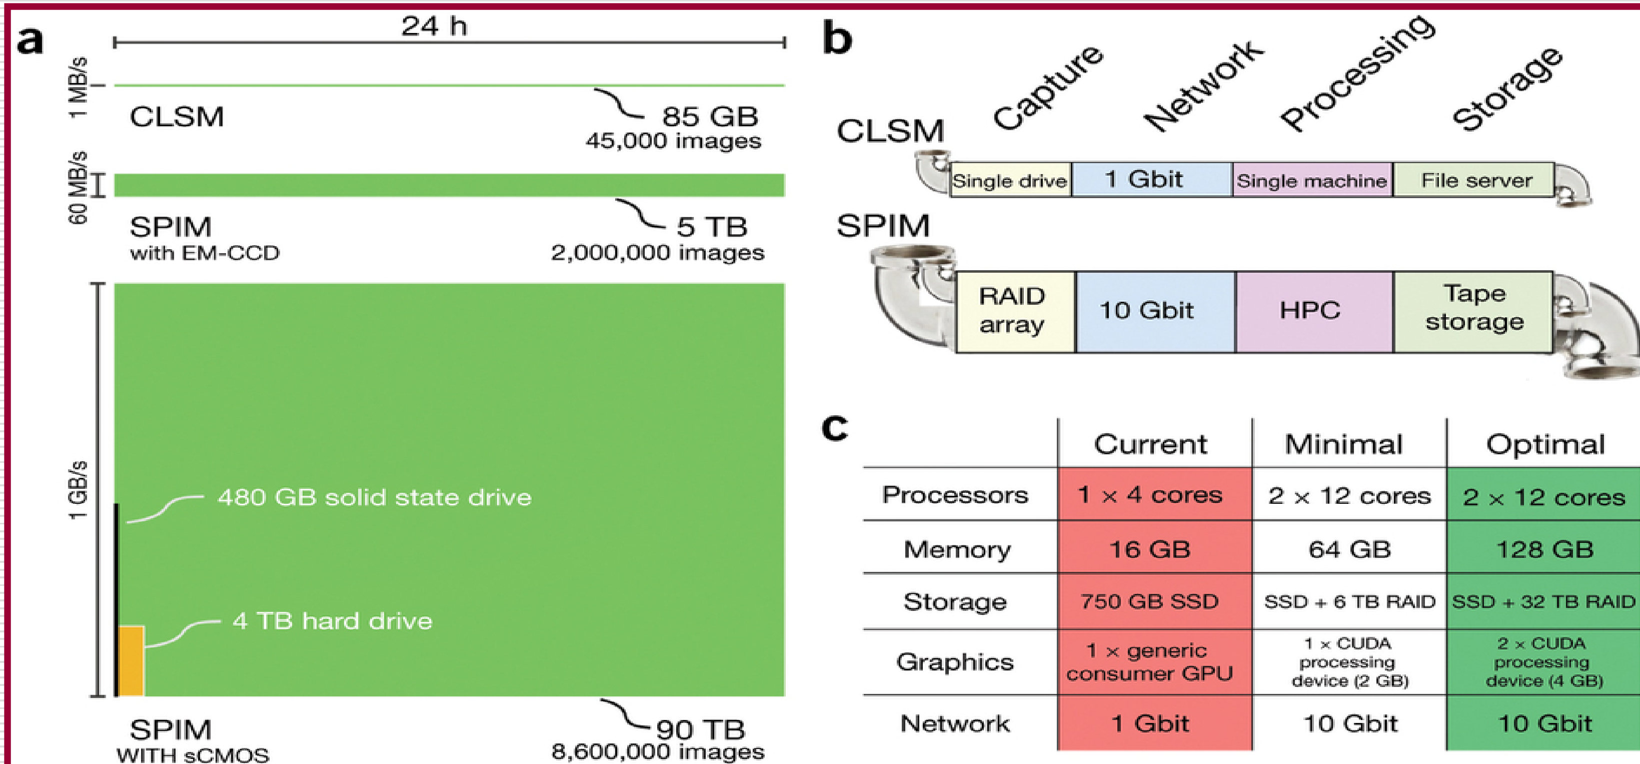

adapted from E. G. Reynaud et al.: **Guide to Lightsheet Microscopy for Adventurous Biologists**, *Nature Methods*, Vol. 12, pp. 30-34, 2015.

New Imaging Technologies:
Lightsheet Microscopy

OpenSPIM (~ **Open** Source **S**elective **P**lane **I**llumination **M**icroscopy)

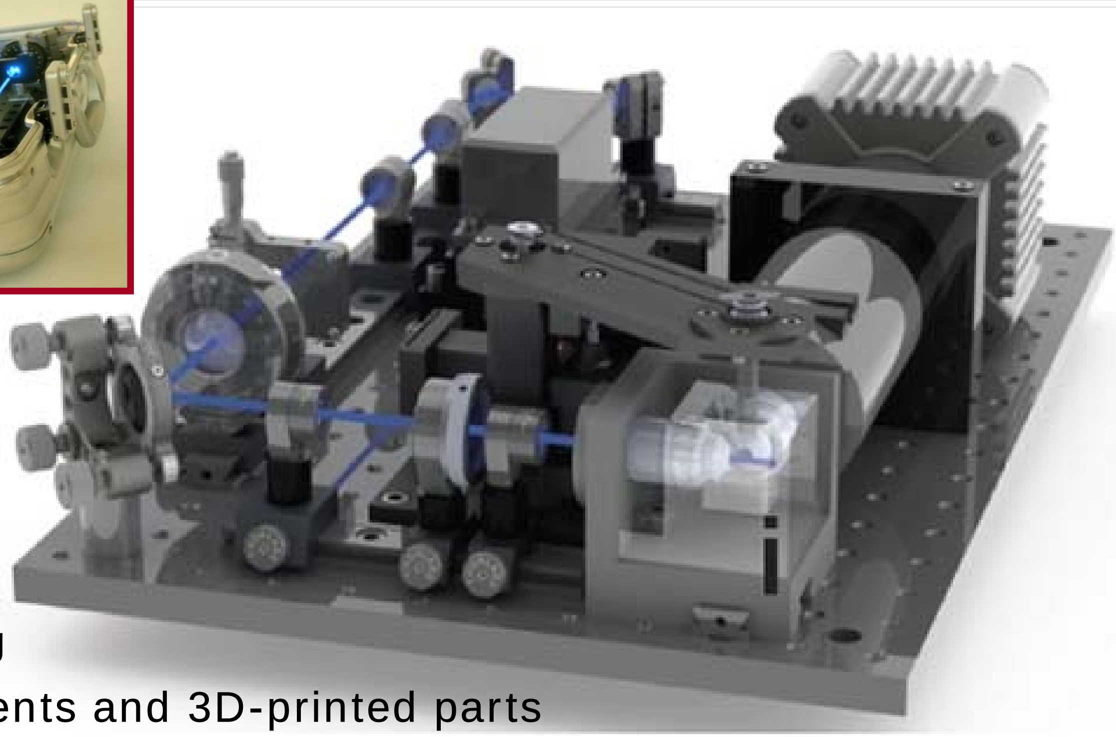

- portable
- cheap
  - ~ 7000 EUR
- easy to assemble
  - http://openspim.org
- off-the-shelf components and 3D-printed parts

Neurobiological Breakthroughs:
Understanding of the visual system

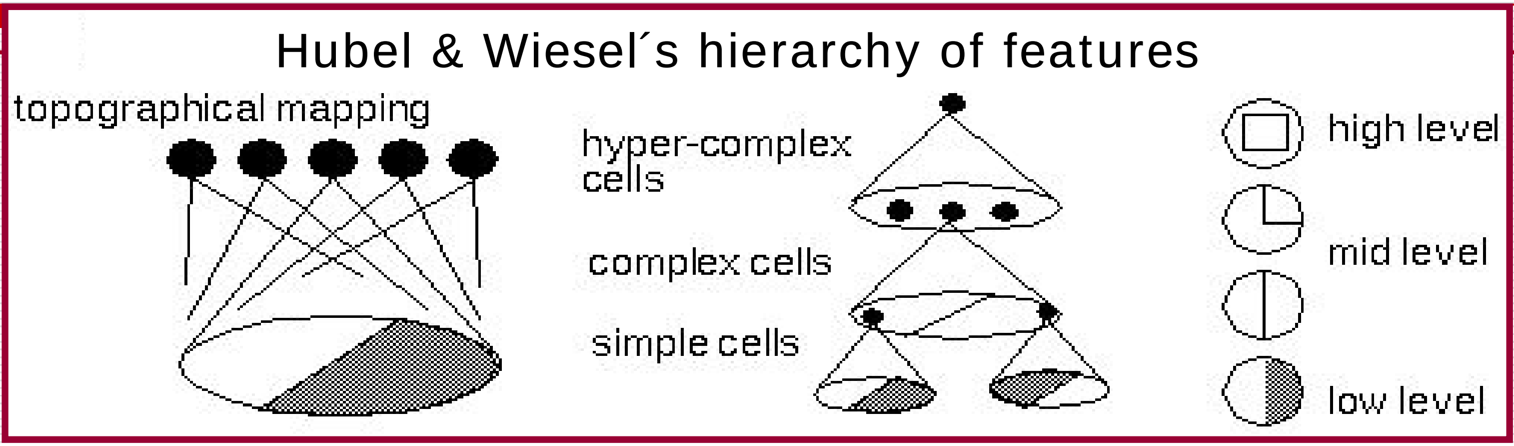

- **based on the research of David H. Hubel and Torsten N. Wiesel on functional architecture in the cat's visual cortex in 1959 and 1962 (Nobel Prize in 1981)**
- **architecture of the visual system** ~ individual cortical cells respond not to the presence of light, but rather to contours of specific orientation; feature-detecting cells form a hierarchy of multiple stages
- **ocular dominance** ~ the preference of cells that process visual stimuli to respond to input from one or the other eye.

  => therapy for children born with cataracts or strabismus

http://cns-alumni.bu.edu/~slehar/webstuff/pcave/hubel.html
https://youtu.be/IOHayh06LJ4

Neurobiological Breakthroughs: Brain Connections

Macaque brain long distance network

https://youtu.be/YZTRxKyx410

Connectograms:

- 2D-graphs of long-distance connections in the brain
- based on in vivo and non-invasively obtained diffusion magnetic resonance imaging data (MRI)
- insight into pathologies

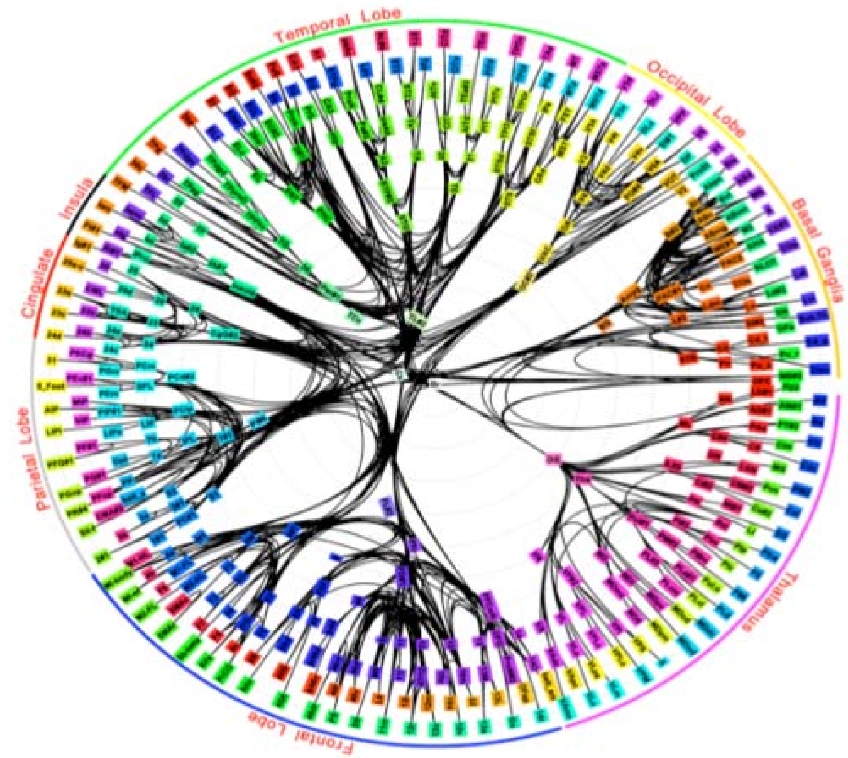

- Dharmendra S. Modha, Raghavendra Singh: Network architecture of the long-distance pathways in the macaque brain, PNAS 2010;107:13485-13490.

Neurobiological Breakthroughs: Brain Connections in Autism

Connectivity between 19 different brain regions, based on EEG data:

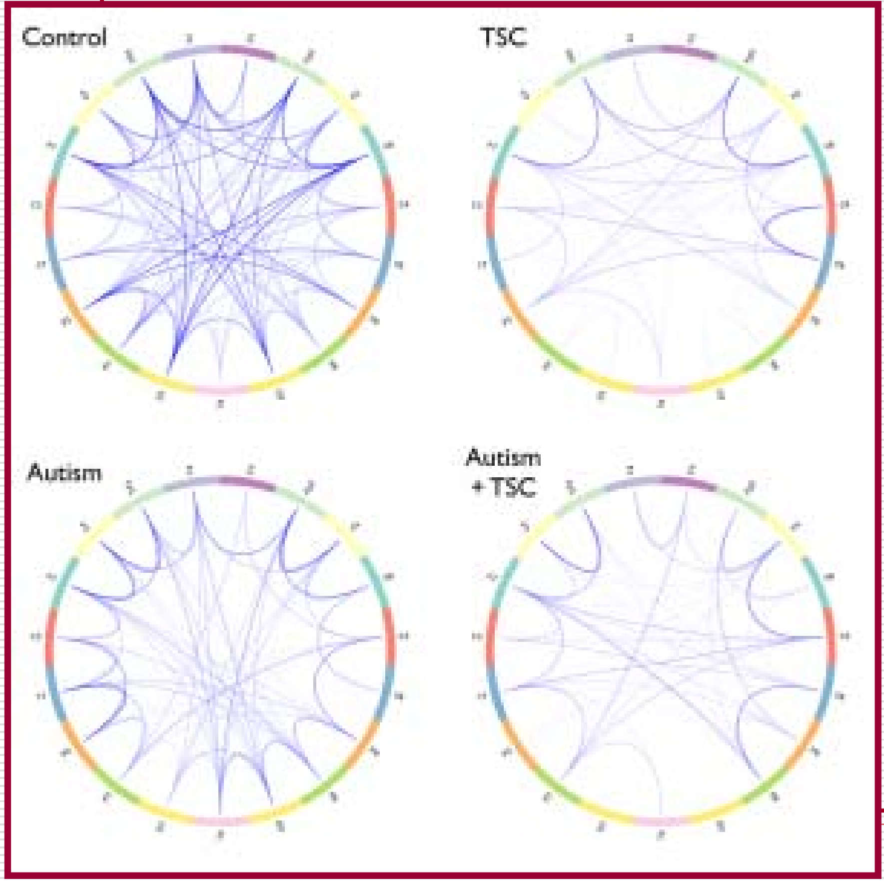

- 46 healthy neurotypical children,
- 16 children with classic autism,
- 14 children whose autism is part of a genetic syndrome called TSC
- 29 children with TSC but not autism

- **Both groups of children with TSC show fewer connections overall**
- **Both groups with autism have more connections between adjacent areas of the brain and fewer connections across distant areas.**

JM Peters et al.:"Brain functional networks in syndromic and non-syndromic autism: a graph theoretical study of EEG connectivity," *BMC Medicine*. Published online Feb. 27 2013

Neurobiological Breakthroughs:

Connections and cortical measures of 110 normal, right-handed males, aged 25-36

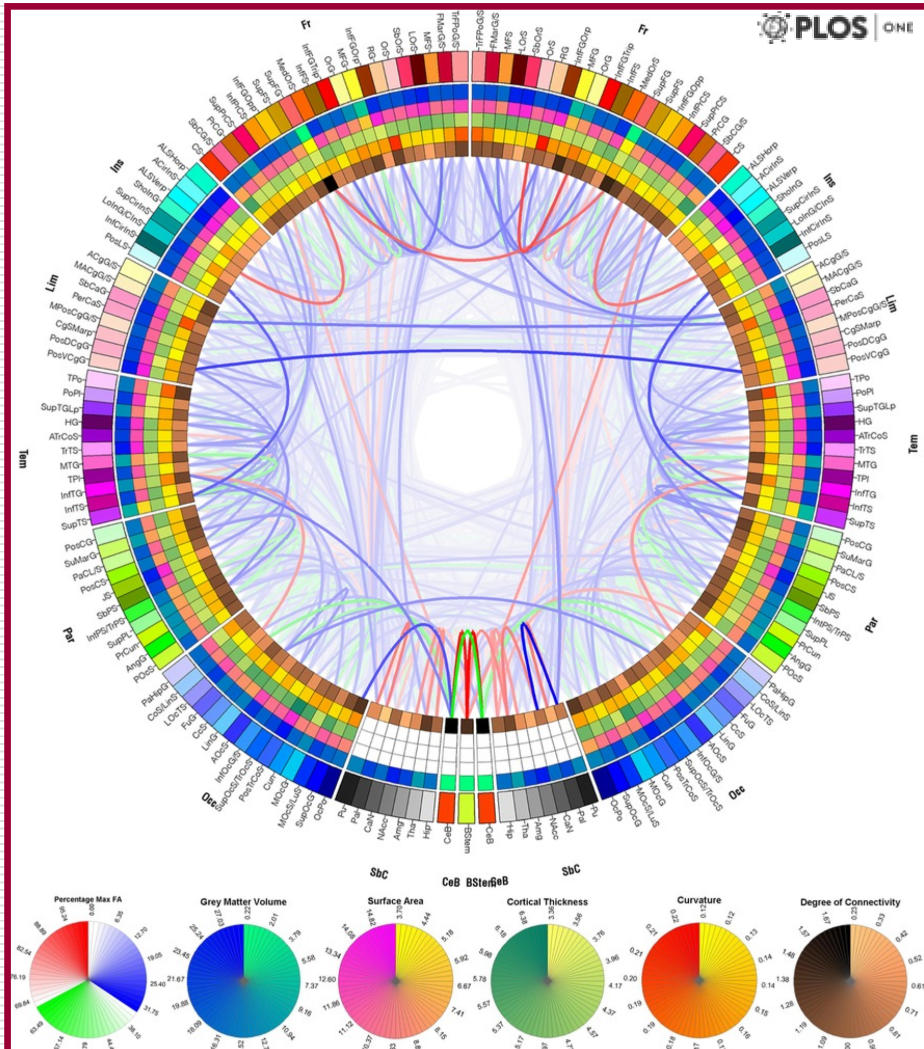

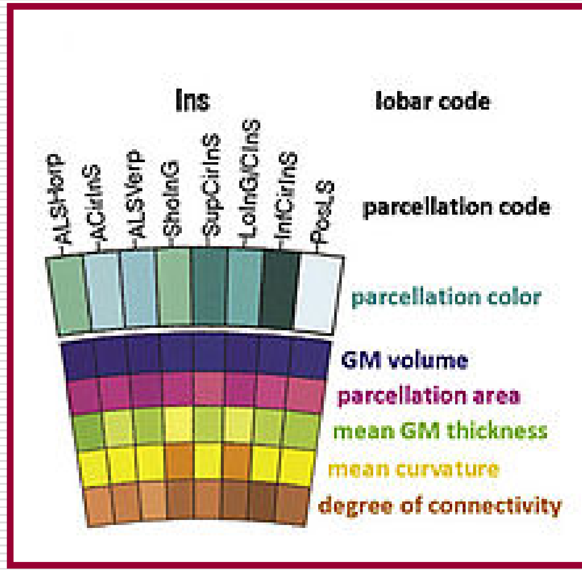

- Van Horn JD, Irimia A, Torgerson CM, Chambers MC, Kikinis R, et al. (2012) **Mapping Connectivity Damage in the Case of Phineas Gage**. PLoS ONE 7(5): e37454. doi:10.1371/journal.pone.0037454

Neurobiological Breakthroughs

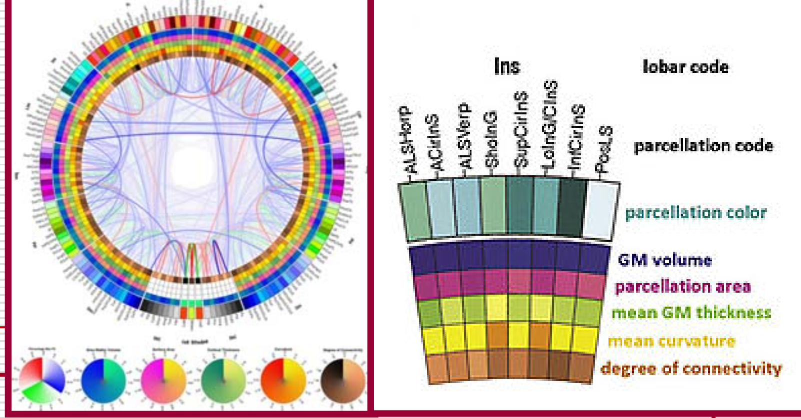

Connectogram with cortical measures:

- 110 normal, right-handed males, aged 25-36
- the left hemisphere is depicted on the left, the right hemisphere on the right
- each cortical area is labeled with an abbreviation and assigned its own color
- the concentric circles represent additional attributes of the corresponding cortical region (grey matter volume, surface area, degree of connectivity, etc.)
- inside the circles, lines connect regions that are structurally connected
- the density (number of fibers) of the connections is reflected in the opacity of the lines

Neurobiological Breakthroughs:
Neuron-Specific Optogenetic Control

Optogenetics ~ brain control with light

- allows for fine manipulation of neuronal activity to control the function of neuronal microcircuits *in vitro* and *in vivo*
- only the genetically targeted cells will be under the control of the light while leaving other cells to function normally
- optical stimulation (light in the UV to the IR wavelengths) can control (either excite or inhibit) genetically targeted neurons in the brain with a high spacial and temporal resolution

Control of social / asocial behavior in mice amygdala

- ChR2 Stimulation of MeApd Neurons Triggers Aggression toward a Female Intruder
- ChR2 Stimulation of vGLUT2$^+$ Neurons Promotes Repetitive Self-Grooming Behavior
- ChR2 Stimulation of vGAT$^+$ Neurons Suppresses Repetitive Self-Grooming Behavior
- http://www.sciencedirect.com/science/article/pii/S0092867414010393

W Hong, D-W Kim, DJ Anderson: Antagonistic Control of Social versus Repetitive Self-Grooming Behaviors by Separable Amygdala Neuronal Subsets, Cell 158 (6), 2014, 1348–1361.

Deep Neural Architectures

- Motivated by biological neural networks
- Some functions compactly represented with k (k>2) layers may require exponential size with 2 layers
- Hierarchy, structure, sparse coding and shared representations
- Various approaches include:
  - Neocognitron
  - Multilayer Perceptrons and Error Back Propagation
  - Convolutional Neural Networks
  - Deep Belief Networks

Neocognitron

- Proposed by Kunihiko Fukushima in 1980
- Kunihiko Fukushima: **Neocognitron: A Hierarchical Neural Network Capable of Visual Pattern Recognition**, *Neural Networks*, Vol. 1, pp. 119-130, 1988.
- **Sparse hierarchical network structure**
- **1D-view of interconnections between the neurons from different layers**

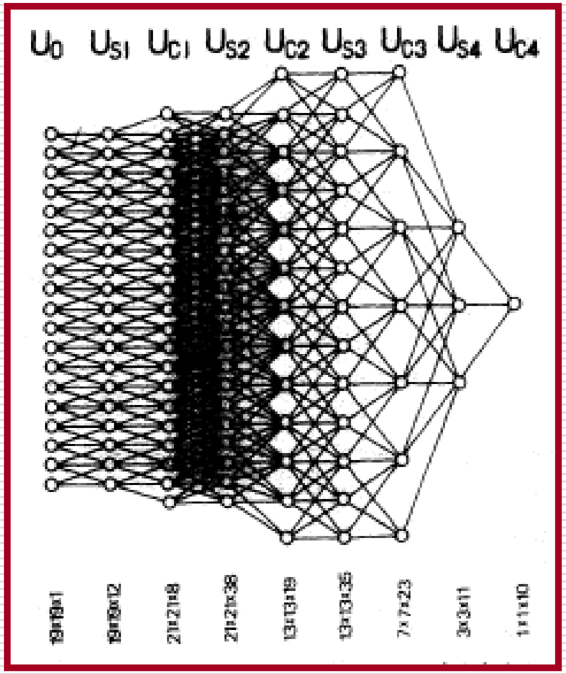

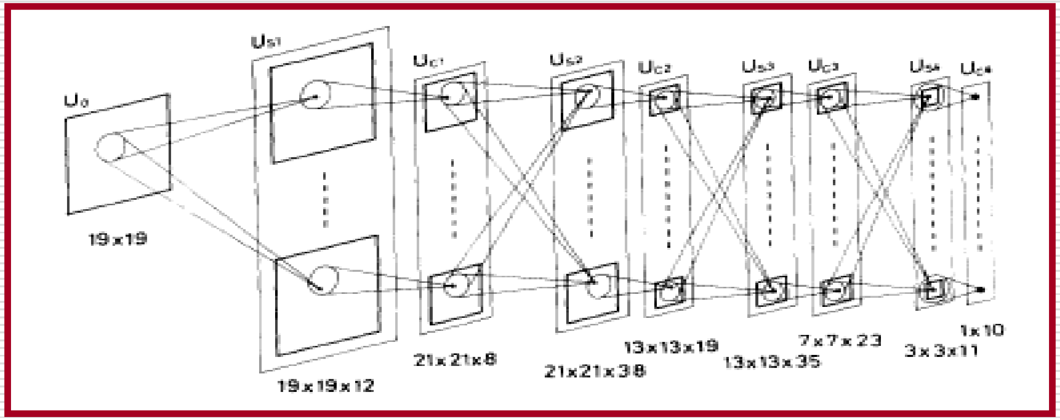

Neocognitron:
two types of neurons

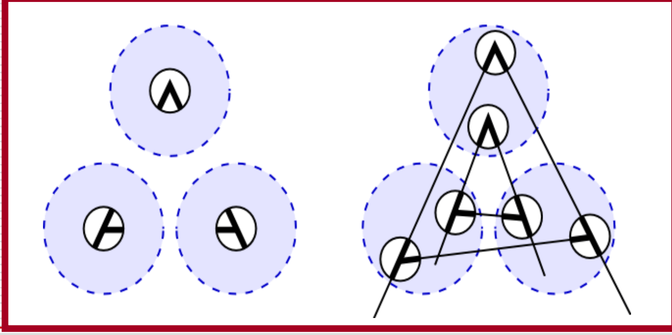

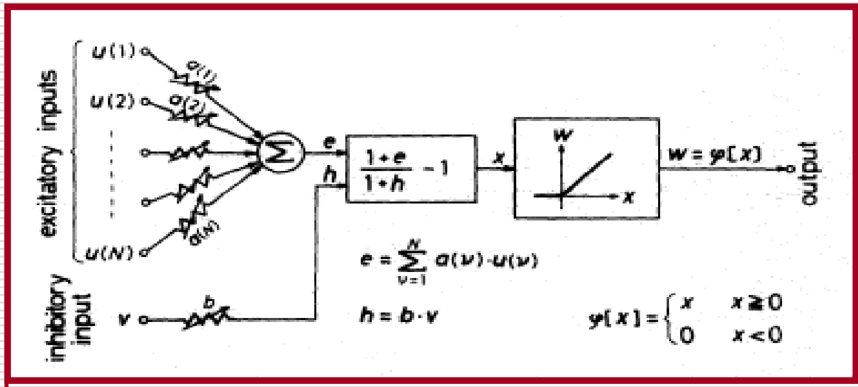

- **S-cells:**
  - extract features at certain positions
  - variable incoming weights reinforced during training
- **C-cells:**
  - support shift invariance in the input
  - fixed incoming weights
  - receive signals from several S-cells extracting the same feature, but at different positions
  - activated if at least some of these S-cell groups is active

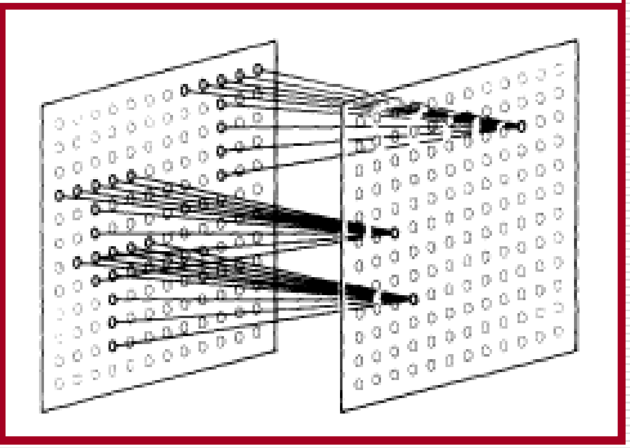

K. Fukushima: **Neocognitron: A Hierarchical Neural Network Capable of Visual Pattern Recognition**, *Neural Networks*, Vol. 1, pp. 119-130, 1988.

Neocognitron:
the recall process

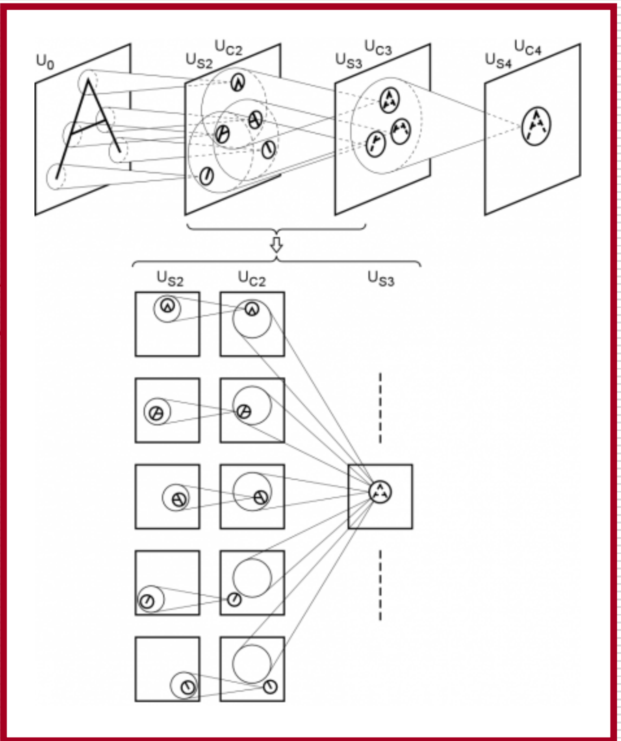

- the cells are arranged into 2D-arrays (~ cell-planes)
- alternating layers of S- and C-cells
- simple features extracted in lower layers are combined into more complex features at higher layers
- the cells at higher layers process larger areas of the input
- neighbouring cells receive similar signals
- at the top layer, there is only 1 C-cell in each cell-plane
  - each of these C-cells is activated only by input patterns from the corresponding category

K. Fukushima: **Neocognitron: A Hierarchical Neural Network Capable of Visual Pattern Recognition**, *Neural Networks*, Vol. 1, pp. 119-130, 1988.

Neocognitron:
the training process

Two main principles:

1. Reinforcement of maximum output cells
   - Only the cell best responding to the training stimulus will be selected to have its weights reinforced
   - Once a cell is selected and its weights reinforced, it usually loses its responsivness to other features
2. Development of iterative connections
   - All the S-cells in the cell-plane respond to the same feature, and the differences between them arise only from the difference in position of the feature to be extracted

Neocognitron - training

1. Initialize the weights with small positive values.
2. Repeat until convergence
   - present an input pattern to the network;
   - in each cell-plane, choose the S-cell with the strongest response (~ the seed cell);
   - reinforce the weights of the input connections for the selected “winning” S-cell to strengthen its response to the detected feature;
   - reinforce also the weights of the input connections for all other S-cells from the same cell-plane using the “winning” cell as a template.

Neocognitron: characteristic properties

Image courtesy
of A. J. Frazer

- A pioneering neural network model capable of learning to recognize 2D-visual patterns
- Robust to errors in position, scale and distortion
- Higher layers can be trained only after the training of preceding stages has been completely finished
- Labeled seed cells are required for supervised training
- During selforganization, maxi-mum output cells are selected automatically as seed cells

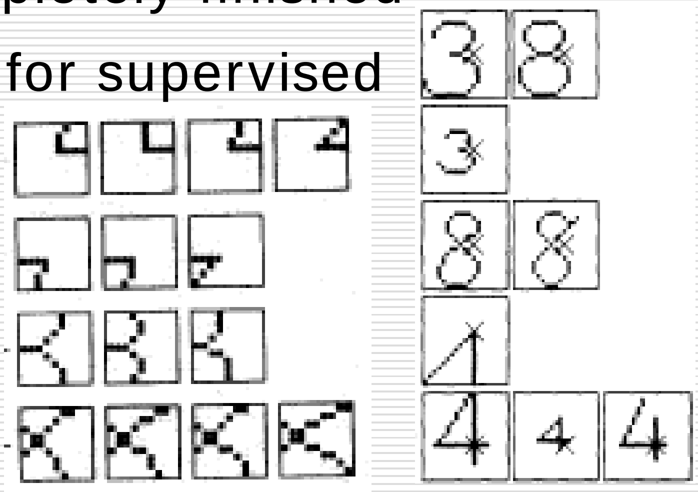

K. Fukushima: **Neocognitron: A Hierarchical Neural Network Capable of Visual Pattern Recognition**, *Neural Networks*, Vol. 1, pp. 119-130, 1988.

Multilayer Perceptrons: Formal Neuron

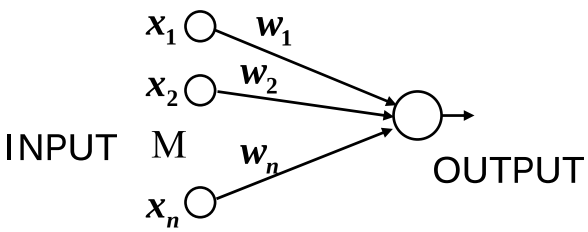

$$\xi = \sum_{i=1}^{n} w_i x_i + \vartheta$$

$$y = f(\xi) = \frac{1}{1 + e^{-\xi}}$$

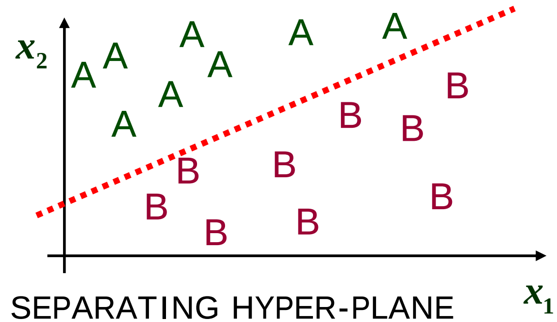

$$x_2 = -\frac{w_1}{w_2} x_1 - \frac{\vartheta}{w_2}$$

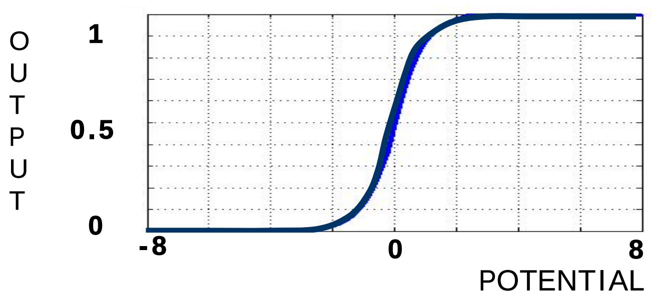

$$y = \begin{cases} 1 \;\; if \; \sum_{i=1}^{n} w_i x_i + \vartheta \geq 0 & : \text{CLASS A} \\ -1 \;\; if \; \sum_{i=1}^{n} w_i x_i + \vartheta < 0 & : \text{CLASS B} \end{cases}$$

Multilayer Perceptrons
and the Error Back Propagation

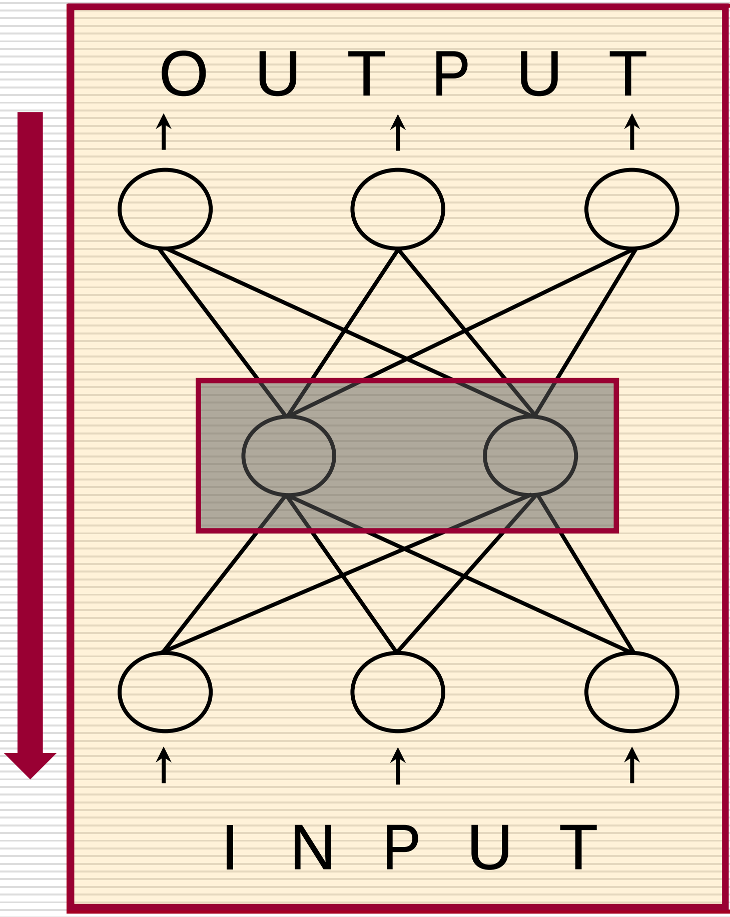

- compute the actual response of the network and compare it with its desired response

$$E = \frac{1}{2} \sum_{p} \sum_{j} \left(y_{j,p} - d_{j,p} \right)^2$$

- **Goal: minimize the error**
  - adjust the weights and thresholds
  - from the output to the input

Multilayer Perceptrons and Error Back Propagation

- First used for gradient evaluation by Paul J. Werbos in 1974

1: **Initialize the weights to small random values**
2: **Present a new training pattern in the form of:** *[input* ***x****, desired output* ***y****]*
3: **Calculate actual output:** in each layer, the activity of neurons is given by:

$$y_j = f(\xi_j) = \frac{1}{1 + e^{-\lambda \xi_j}}, \qquad \text{where} \qquad \xi_j = \sum_i y_i w_{ij}$$

4: **Weight adjustment:** start at the output layer and proceed back towards the input layer according to:

$$w_{ij}(t+1) = w_{ij}(t) + \alpha \delta_j y_i + \alpha_m \left(w_{ij}(t) - w_{ij}(t-1)\right)$$

$$\delta_j = \begin{cases} (d_j - y_j)\,\lambda\, y_j(1 - y_j) & \text{for an output neuron} \\ \left(\sum_k \delta_k w_{jk}\right)\lambda\, y_j(1 - y_j) & \text{for a hidden neuron} \end{cases}$$

for the weight $w_{ij}(t)$ from neuron i to neuron j in time t; learning/momentum rates α / α_m ; potential / local error on neuron j denoted as ξ_j / δ_j ; the index k for the neurons from the layer above the neuron j and the slope of the transfer function λ

5: **Repeat by going to step 2**

D. E. Rumelhart, G. E. Hinton, R. J. Williams: **Learning Representations by Back-Propagating Errors**, *Nature*, Vol. 323, pp. 533-536, 1986.

George´s Girls

Task: guess if George will like that girl

	hair length	intelligence	sense of humor	blue eyes	1.hidden neuron	2. hidden neuron	**attractivity**
1.	0.84	0.39	0.78	0.79	0.64	1.00	**0.42**
2.	0.91	0.19	0.33	0.77	0.00	1.00	**0.20**
3.	0.27	0.55	0.47	0.69	0.98	1.00	**0.50**
4.	0.36	0.51	0.95	0.91	0.86	1.00	**0.60**
5.	0.63	0.71	0.14	0.61	0.85	1.00	**0.62**
6.	0.02	0.24	0.13	0.80	0.02	1.00	**0.05**
7.	0.61	0.69	0.63	0.52	1.00	1.00	**0.80**
8.	0.49	0.97	0.29	0.77	0.59	1.00	**0.40**

Multilayer Perceptrons:
what are the neurons really doing?

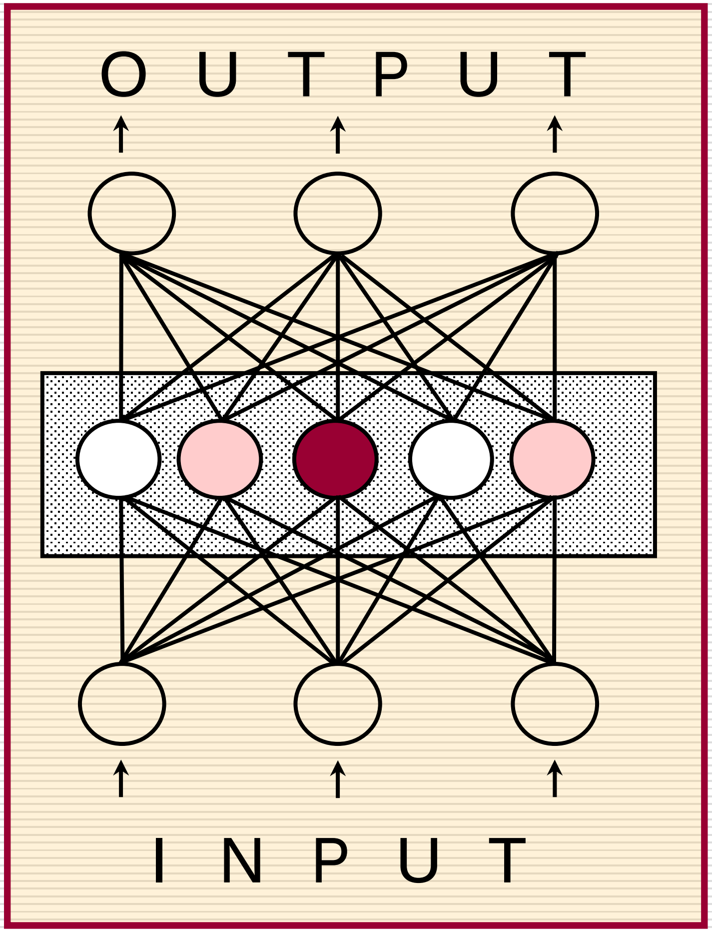

- activity interpretation for hidden neurons:

1	⟷	active	⟷	YES
0	⟷	passive	⟷	NO
0.5	⟷	silent	⟷	DON´T KNOW

- transparent structure
- detection and pruning of redundant neurons
- improved generalization

George´s girls revisited

How many neurons will George need to solve his problem?

	hair length	intelligence	sense of humor	blue eyes	1.hidden neuron	2. hidden neuron	**attractivity**
1.	0.84	0.39	0.78	0.79	0.64	1.00	**0.42**
2.	0.91	0.19	0.33	0.77	0.00	1.00	**0.20**
3.	0.27	0.55	0.47	0.69	0.98	1.00	**0.50**
4.	0.36	0.51	0.95	0.91	0.86	1.00	**0.60**
5.	0.63	0.71	0.14	0.61	0.85	1.00	**0.62**
6.	0.02	0.24	0.13	0.80	0.02	1.00	**0.05**
7.	0.61	0.69	0.63	0.52	1.00	1.00	**0.80**
8.	0.49	0.97	0.29	0.77	0.59	1.00	**0.40**

The German Traffic Sign Competition (IJCNN 2011)

- **Convolutional Neural Networks performed best!**
  - No need for custom-made image pre-processing
  - 98.98 % (Schmidhuber et al), 98.97 % (LeCun et al), 98.81 % (human performance)

CNN-networks
(Convolutional neural networks)

The LeNet-5 model (Yan LeCun et al. 1998)

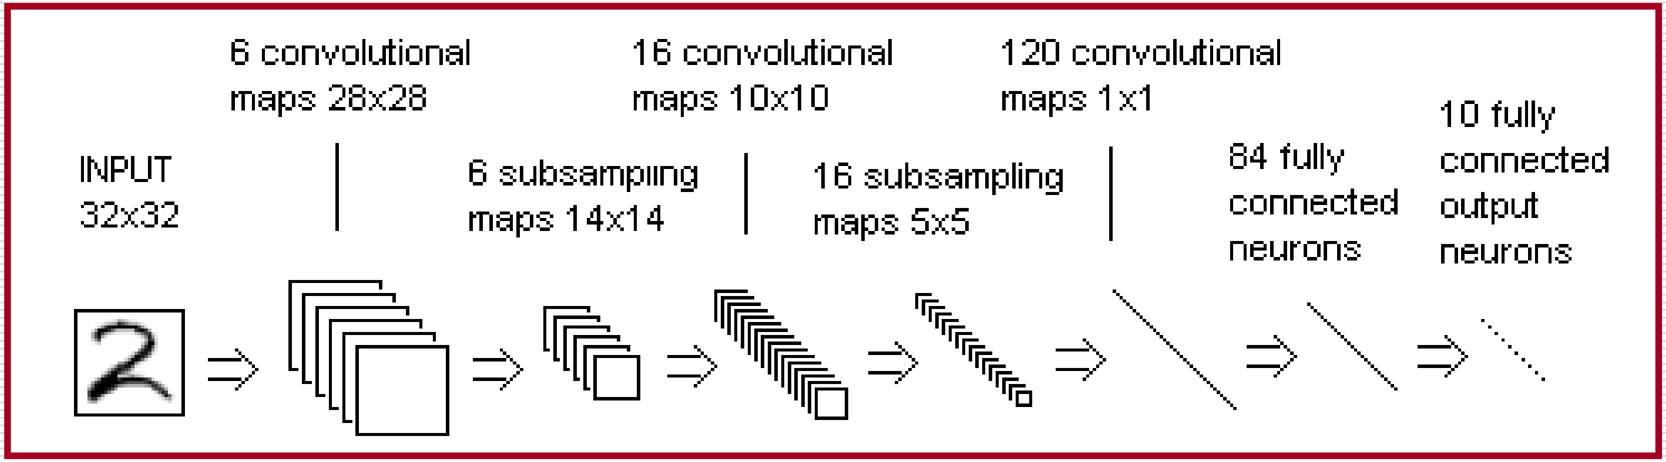

- Trained by back-propagation (sparse connectivity)
- Local receptive fields, weight sharing and spatial sub-sampling (alternating convolutional and subsampling layers)
- Invariant object recognition (up to a certain degree)
- X **Fixed number of feature maps in each layer!**

Y. LeCun, L. Bottou, Y. Bengio, P. Haffner: **Gradient-Based Learning Applied to Document Recognition**, *Proc. of the IEEE*, Vol. 86, pp. 2278–2399, 1998.

CNN-networks:
the convolutional layer *l*

- receptive fields of the same size $r_c^l \times r_c^l$ overlapping in $r_c^l - 1$ rows/columns
- neuron *(i,j,f,l)* at the position *(i,j)* in the feature map *f* of the layer *l* is thus connected to neurons *(i+Δi, j+Δj, f´, l-1)* from the layer *l-1* by the weight $w_{\Delta i,\Delta j}^{f,f´,l}$ for $f´ \in F_{f,l}^{IN}, 0 \le \Delta i, \Delta j \le r_c^l + 1$; the neurons from the feature map *f* take their input from a set of feature maps $F_{f,l}^{IN} \neq \{\ \}$
- The weights are shared for all the neurons from the same feature map
- **The potential $\xi_{i,j}^{f,l}$ and output $y_{i,j}^{f,l}$ of the neuron *(i,j,f,l)*:**

$$y_{i,j}^{f,l} = \xi_{i,j}^{f,l} = \sum_{f´ \in F_{f,l}^{IN}} \sum_{\Delta i=0}^{r_c^l - 1} \sum_{\Delta j=0}^{r_c^l - 1} y_{i+\Delta i, j+\Delta j}^{f´,l-1} \cdot w_{\Delta i,\Delta j}^{f,f´,l}$$

- The size m_l x n_l of all feature maps from *l* is imposed by the size of the feature maps in layer *l-1* and by the size of the receptive field

CNN-networks:
the subsampling layer *l*

- non-overlapping subsampling areas of the size $r_x^l \times r_y^l$ (usually 2x2)
- multiplicative trainable coefficients $a^{f,l}$ and additive trainable biases $b^{f,l}$
- **The potential $\xi_{i,j}^{f,l}$ and output $y_{i,j}^{f,l}$ of the neuron *(i,j,f,l)*:**

a) averaging:
$$\xi_{i,j}^{f,l} = \frac{1}{r_x^l \cdot r_y^l} \sum_{\Delta i=0}^{r_x^l-1} \sum_{\Delta j=0}^{r_y^l-1} y_{i\cdot r_x^l+\Delta i,\, j\cdot r_y^l+\Delta j}^{f',l-1}$$

b) maximizing:
$$\xi_{i,j}^{f,l} = \max_{\substack{\Delta i\in\{0,\ldots,r_x^l-1\}\\ \Delta j\in\{0,\ldots,r_y^l-1\}}} \left(y_{i\cdot r_x^l+\Delta i,\, j\cdot r_y^l+\Delta j}^{f',l-1} \right)$$

$$y_{i,j}^{f,l} = f\left(a^{f,l} \cdot \xi_{i,j}^{f,l} + b^{f,l}\right)$$

- The size m_l x n_l of the feature maps from *l* : $m_l = m_{l-1}/r_x^l$, $n_l = n_{l-1}/r_y^l$

DBN-networks

(Deep Belief Networks) G.E. Hinton et al. 2006

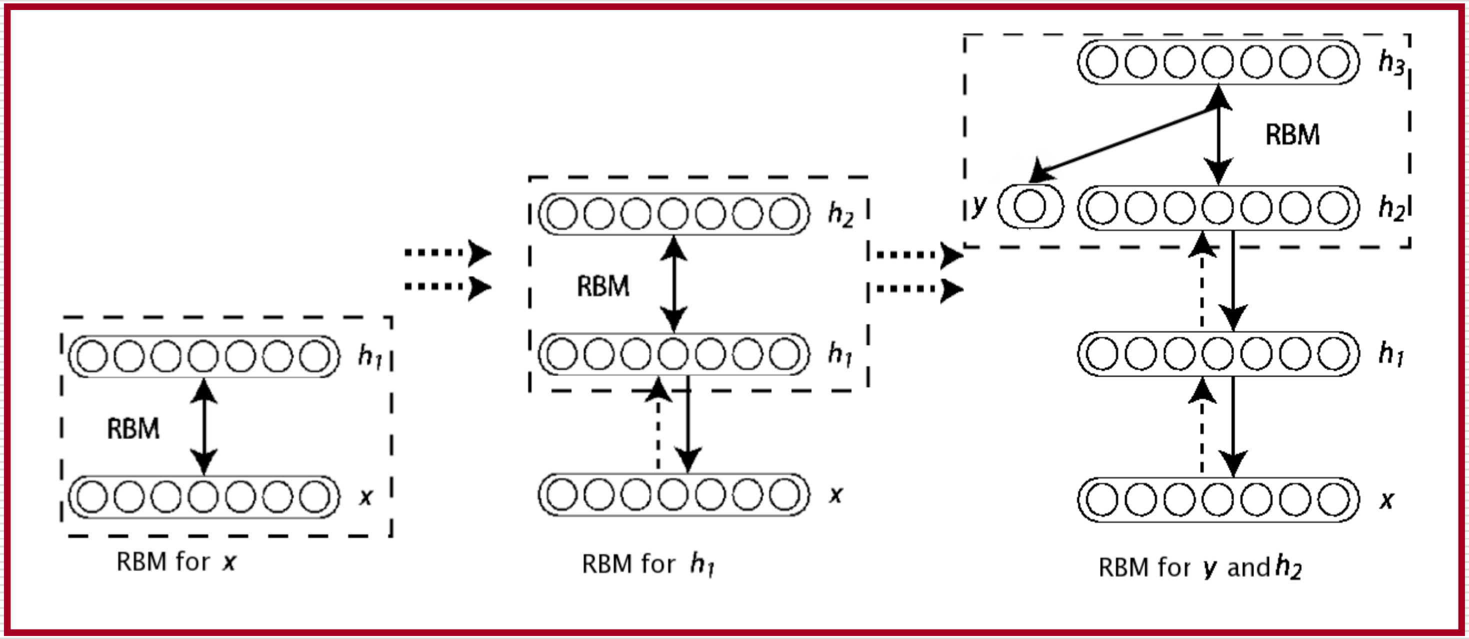

- Stacked Restricted Boltzmann Machines with a classifier
- Unsupervised pre-training (layer-wise)
- Short supervised fine-tuning

G. E. Hinton, S. Osindero, Y.-W. Teh: **A Fast Learning Algorithm for Deep Belief Nets**, *Neural Computation*, Vol. 18, pp. 1527–1554, 2006.

RBM-networks
(Restricted Boltzmann Machines)

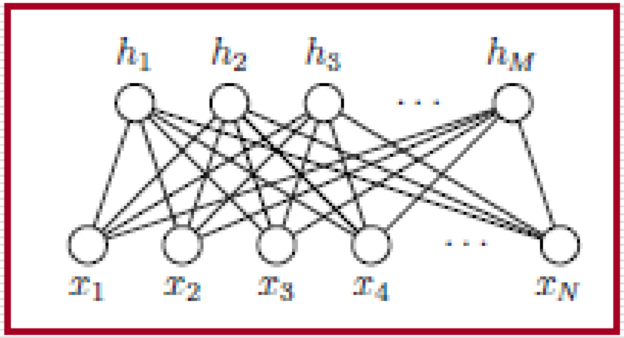

hidden layer: $\mathbf{h} = (h_1, h_2, \ldots, h_M)$

visible layer: $\mathbf{x} = (x_1, x_2, \ldots, x_N)$

$$p\left(h_j = 1 \middle| \mathbf{x}\right) = \frac{1}{1 + e^{-\sum_i w_{ij} x_i - b_j}} \qquad p\left(x_i = 1 \middle| \mathbf{h}\right) = \frac{1}{1 + e^{-\sum_j w_{ij} h_j - c_i}}$$

- A popular building block for deep architectures
- A bipartite undirected graphical model
- RBMs are universal approximators (with enough hidden units, they can perfectly model any discrete distribution)
- Adding one hidden unit (with a proper choice of parameters) guarantees increasing likelihood

N. Le Roux, Y. Bengio: Representational power of restricted Boltzmann machines and deep belief networks, *Neural Computation*, Vol. 20(6) pp. 1631–1649, 2008.

RBM-networks

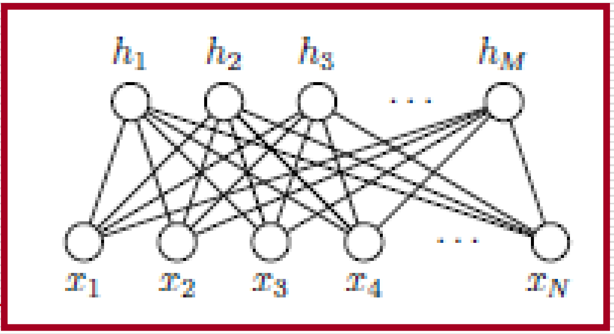

- **Energy function:** $E(\mathbf{x},\mathbf{h}) = -(\mathbf{x}^T\mathbf{W}\mathbf{h} + \mathbf{b}^T\mathbf{h} + \mathbf{c}^T\mathbf{x})$
- Probability of configuration $(\mathbf{x},\mathbf{h})$: $p(\mathbf{x},\mathbf{h}) = \dfrac{e^{-E(\mathbf{x},\mathbf{h})}}{\sum_{\mathbf{x}',\mathbf{h}'} e^{-E(\mathbf{x}',\mathbf{h}')}}$
- **Our goal:** $p(\mathbf{x}) = p_{train}(\mathbf{x})$

  ==> maximize the likelihood of the training data
- As $\dfrac{\partial \log p(x)}{\partial w_{ij}} = x_i^0 h_j^0 - x_i^\infty h_j^\infty \approx x_i^0 h_j^0 - x_i^k h_j^k$

  adjust the weights by: $w_{ij}^{t+1} = w_{ij}^t + \alpha \dfrac{\partial \log p(x)}{\partial w_{ij}}$

  (and similarly for the biases)

Applications

- Image Classification / Processing
- Signal and Multimedia Data Processing
- Knowledge Extraction and Interpretation

Recognition of handwritten digits

accuracy of CNN-networks around 93 % (with M. Kukacka)

- Simple local primitives: e.g., **background**, **background followed by an object from the right**, **diagonal line**

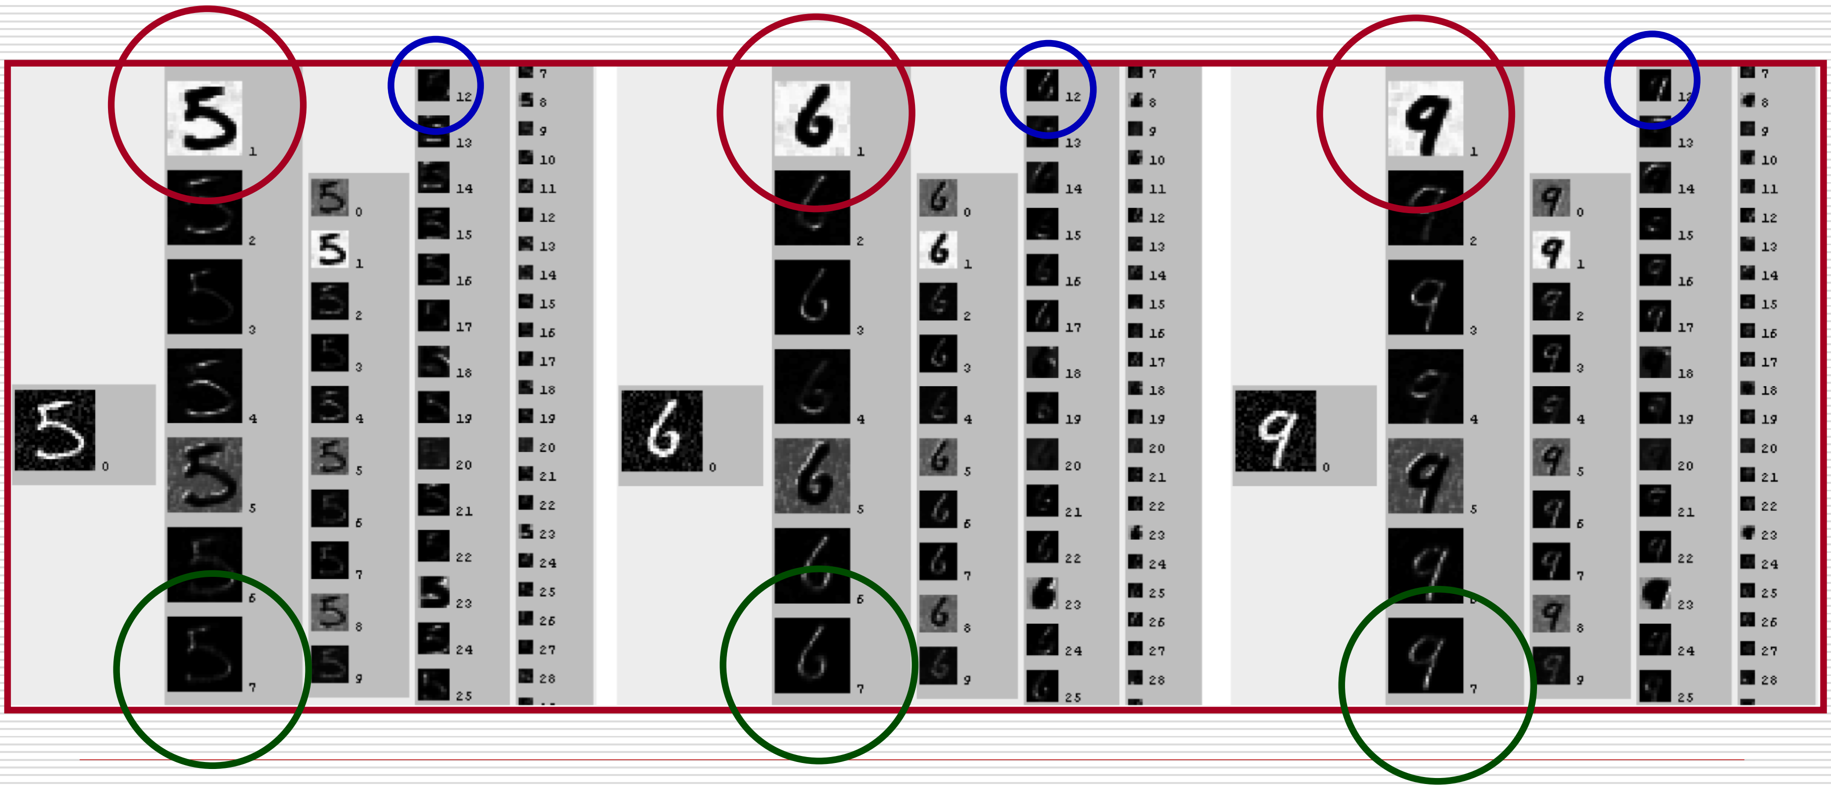

Recognition of human faces

accuracy of CNN-networks almost 93 % (with M. Kukacka)

Features:

- **light surfaces**
- **dark surfaces**
- **light-coloured noses**

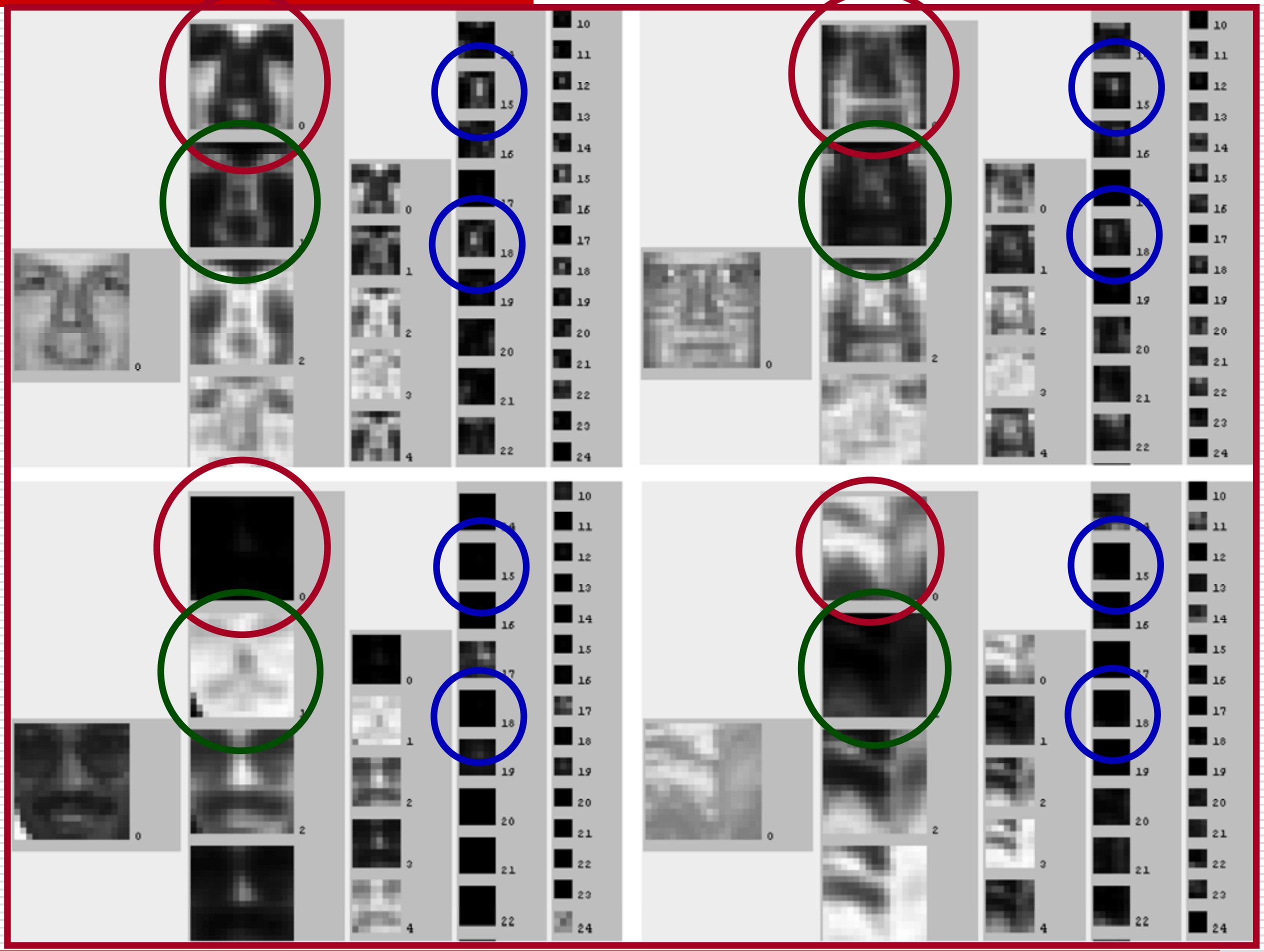

Detection of Hockey Players

accuracy of CNN-networks over 98.5 % (with M. Hrincar)

Objective:

- Reliable online video processing for augmented reality

Data:

- Records of broadcasted hockey matches (publicly available during the World Championships 2011 and 2012)

Results:

- http:tinyurl.com/hokejdetect

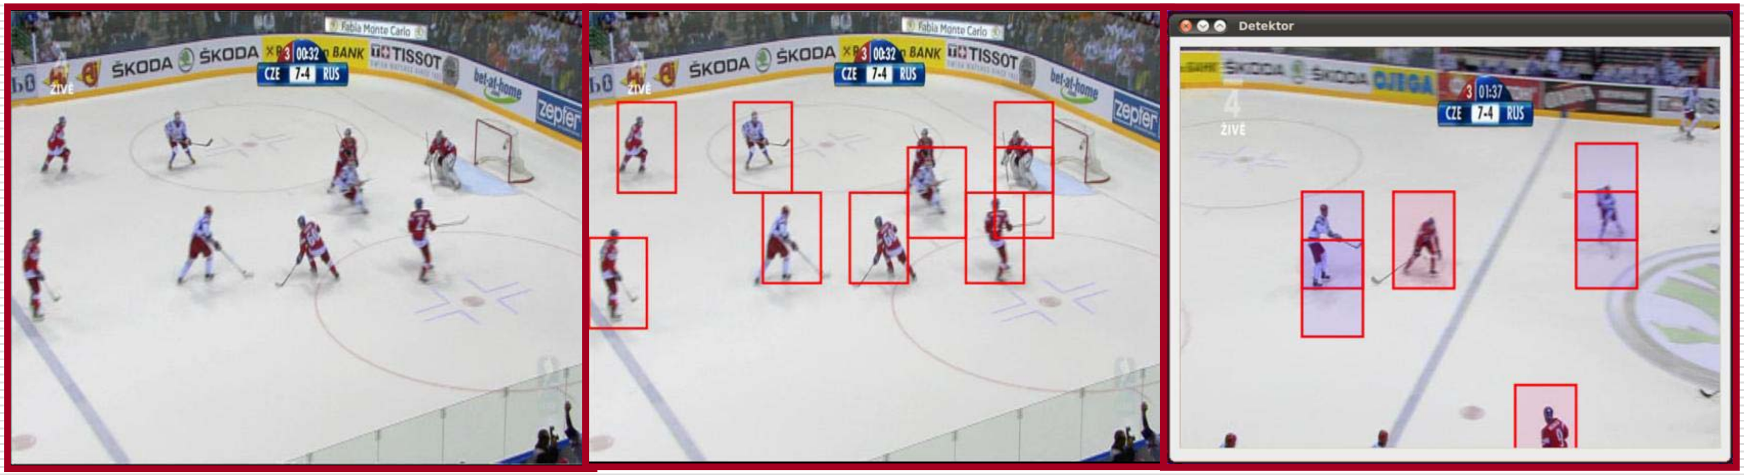

Detection of Hockey Players

accuracy of CNN-networks over 98.5 % (with M. Hrincar)

- accuracy of a CNN-network trained on original data to Gaussian noise (with zero mean and growing variance)

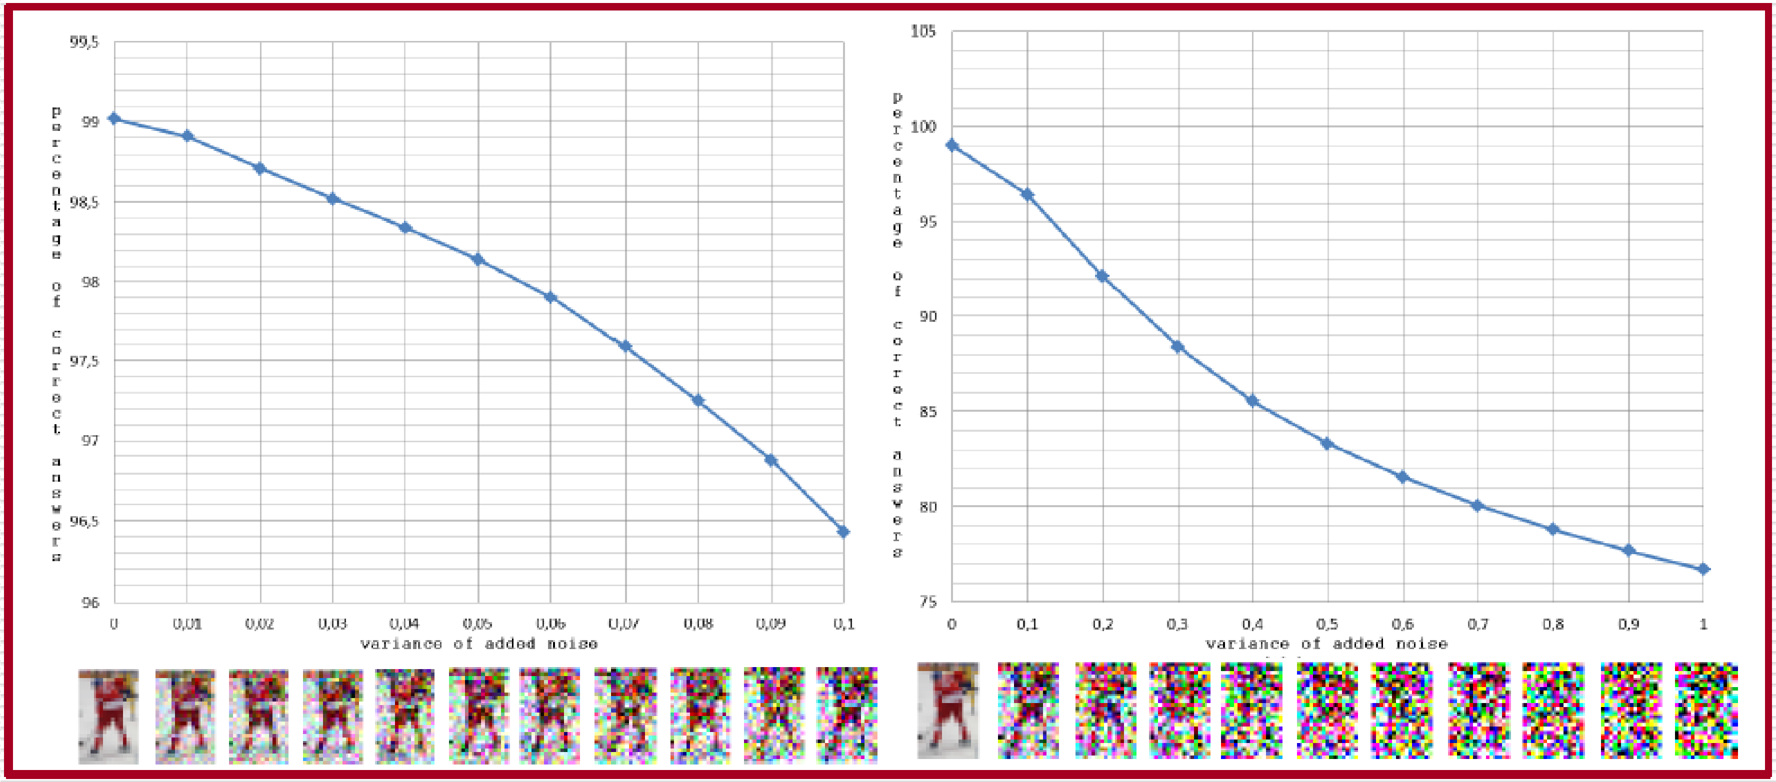

Detection of Hockey Players

accuracy of CNN-networks over 98.5 % (with M. Hrincar)

- internal representations in the feature maps filter out the noise

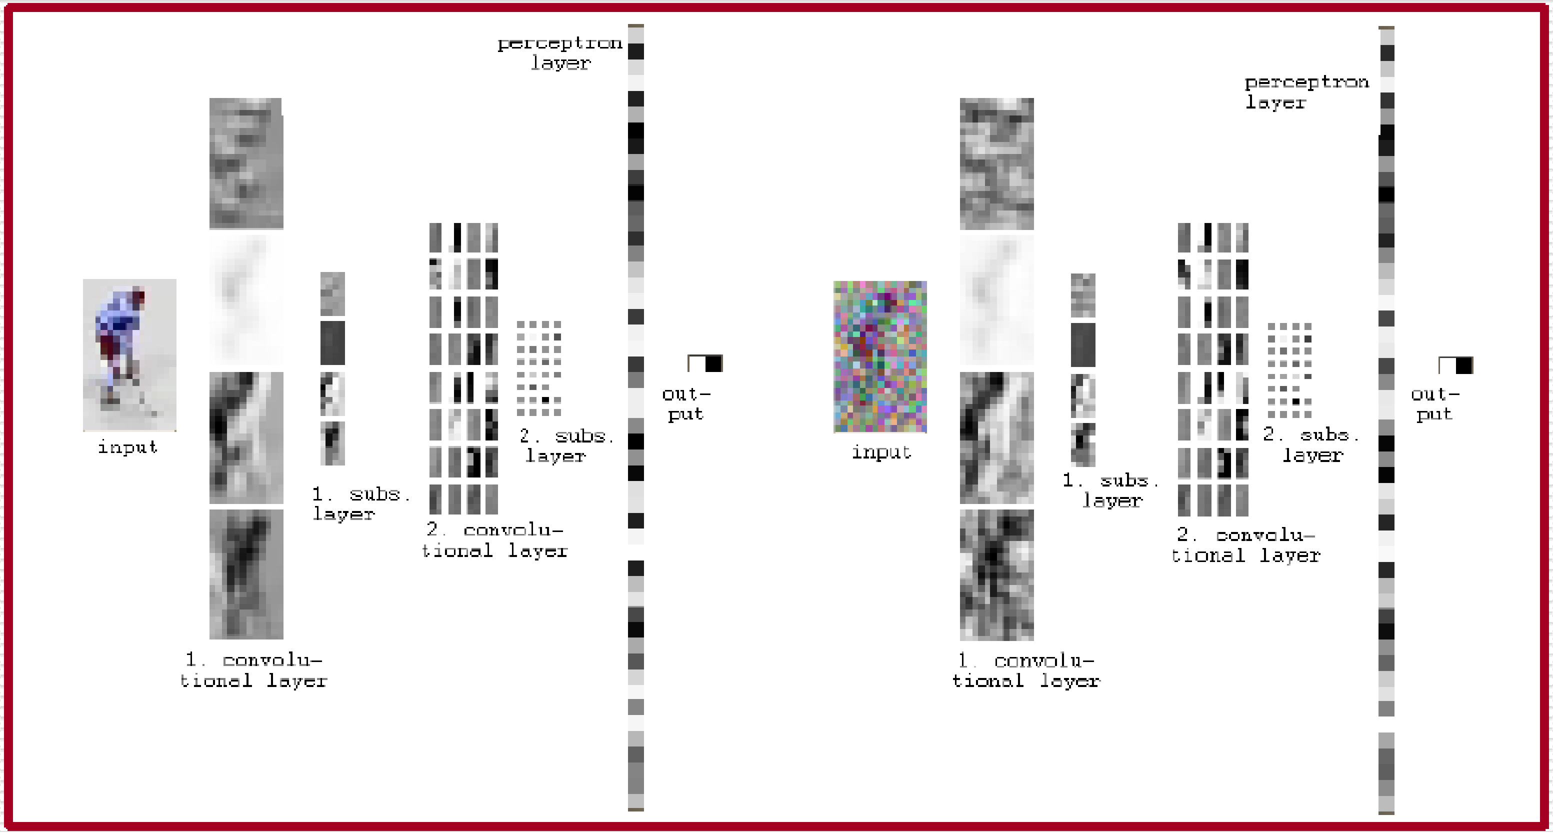

Deep Neural Networks for 3D-data Processing (with J. Pihera and J. Veleminska)

- Detection of characteristic face features
- Classification of 3D-face models according to the person´s gender

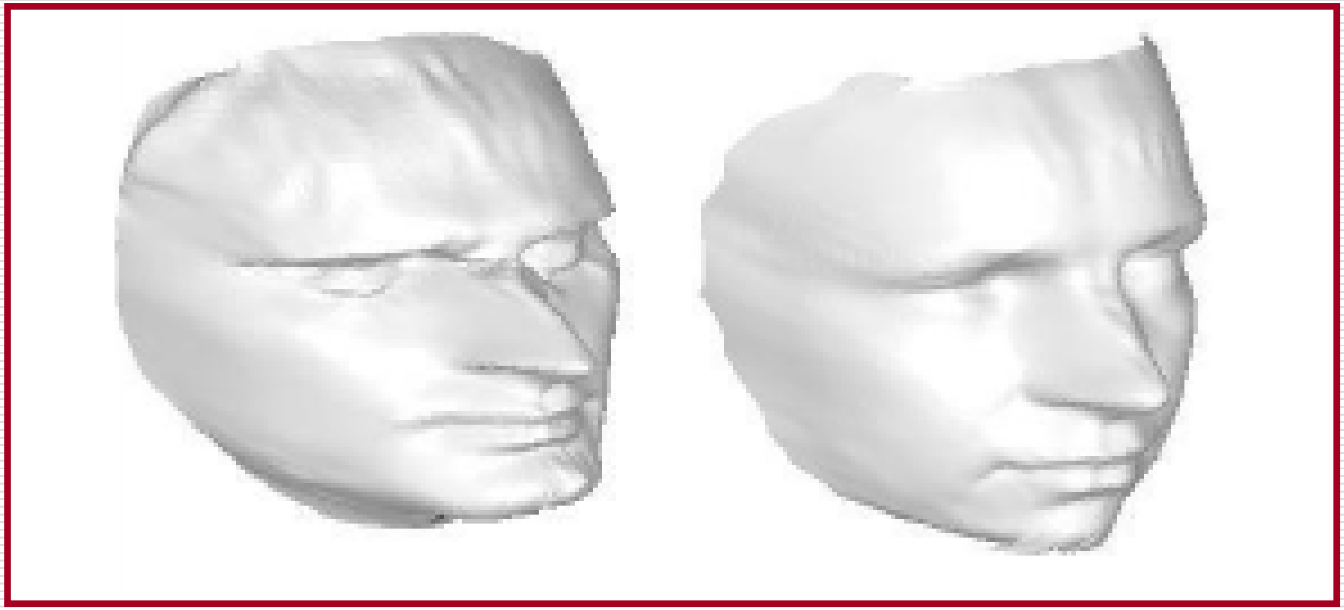

George´s Girls – Are That Girls?

Model - Andrej Pejic
(source: idnes.cz)

Transsexual participant of Miss Universe Canada
(source: idnes.cz)

Miss Tiffany’s Universe trans-genders
(source: super.cz)

- **Difficult to determine gender based on the face**

Human performance (accuracy) on 3D-face scan classification:

Gender of the participants	Gender of the faces		
	Men	Women	All
Men	78.1 %	54.5 %	65.2 %
Women	76.8 %	51.0 %	62.6 %
All	77.4 %	52.7 %	63.8 %

Data and models

- 3D data – face models

  (courtesy of the Department of Anthropology and Human Genetics, Faculty of Natural Sciences of the Charles University in Prague)

- Theoretical model
  - Kohonen´s SOM
  - GNG (Growing Neural Gas)
  - Convolutional Neural Networks ~ an advanced model for shape recognition in 2D-images

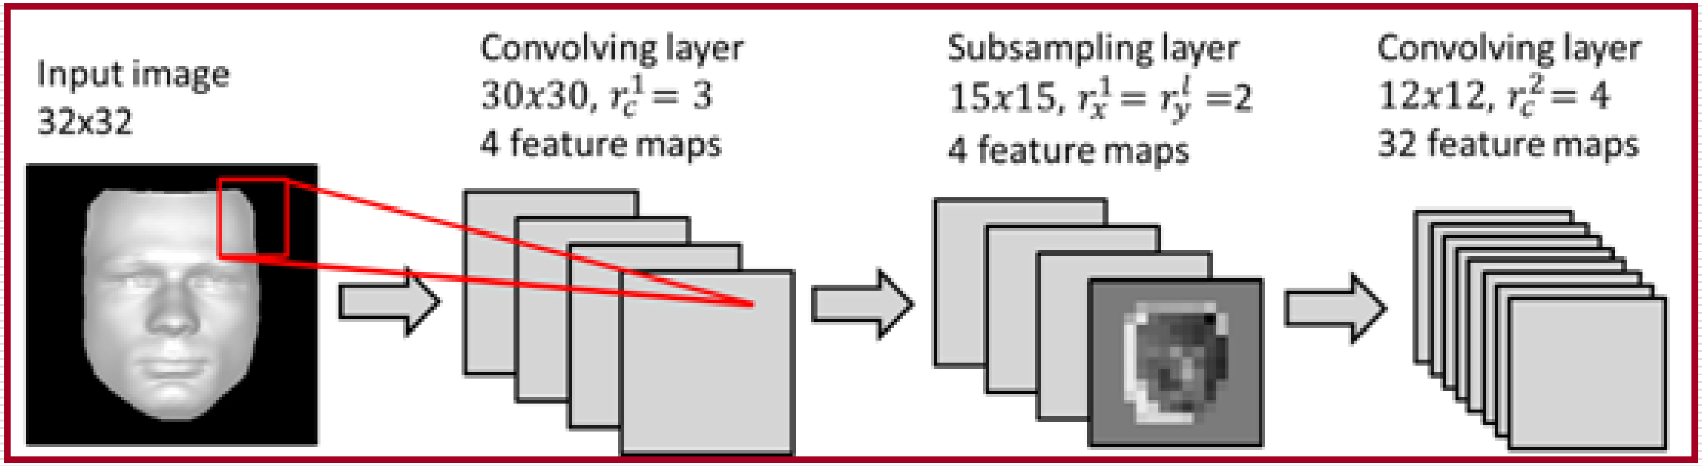

Detection of characteristic face features (with J. Pihera and J. Veleminska)

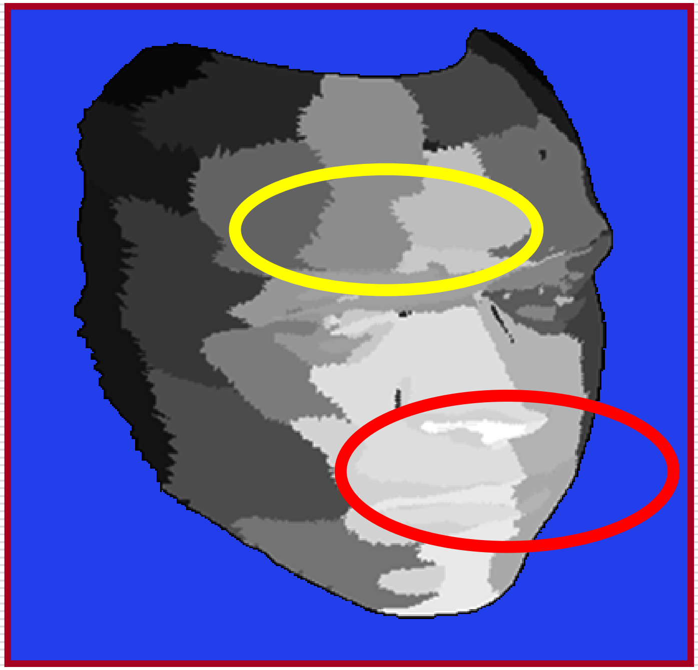

SOM, 20x20 neurons, 34 clusters

GNG, 400 neurons, 40 clusters

– self-organizing neural network models trained on the face data
– clustering of the neurons and labeling of the clusters

Sexual Dimorphism – classification according to person´s gender

(with J. Pihera and J. Veleminska)

a 2D-transform:

- a drawn 3D model (raw)
- pre-processing by means of a SOM
- images

a 3D transform:

- direct / clustered
- pre-processing by means of a SOM
- 3D tensors (22^3 voxels)

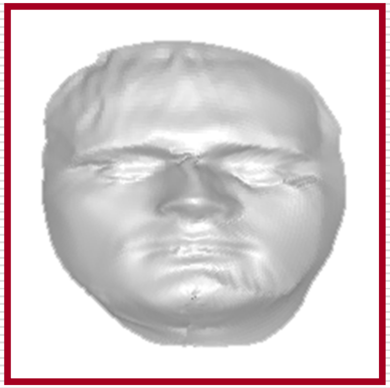

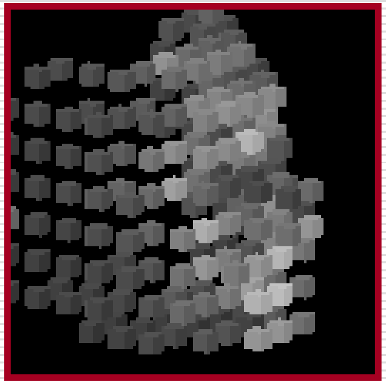
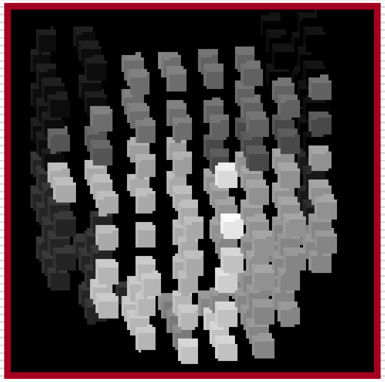

Examples of rotated and scaled patterns added to the training set.

Classification according to person´s gender – 2D

(with J. Pihera and J. Veleminska)

Input and output of the first detection layer

Raw:

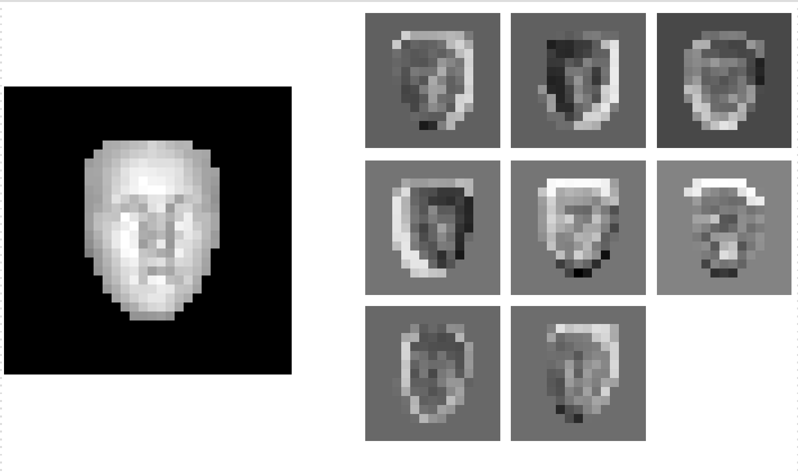

Pre-processed by a SOM:

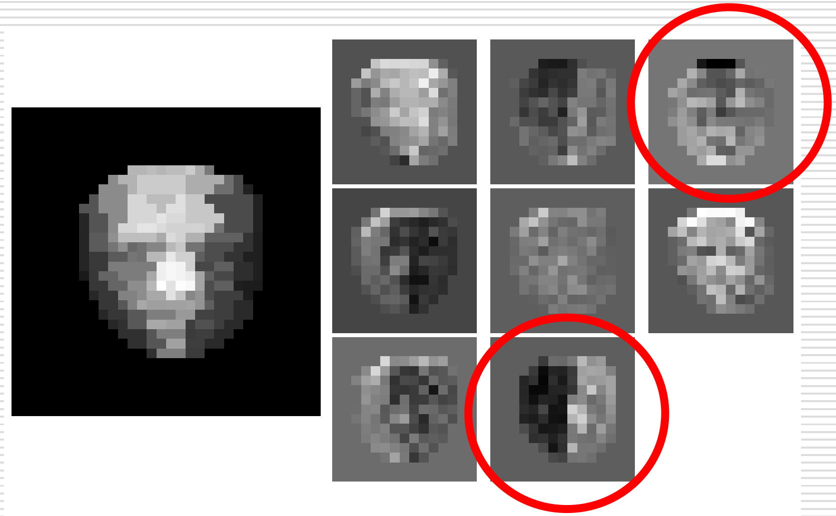

Classification according to person´s gender – 3D

(with J. Pihera and J. Veleminska)

- Convolutional neural networks were designed to process 2D-information
- 3D tensors at the input
- New model of ND-CNNs:
  - Extend the feature maps to process N-dimensional object information
- Feature maps shrink very fast *x* combine the input from a large region
- Complexity similar to CNNs

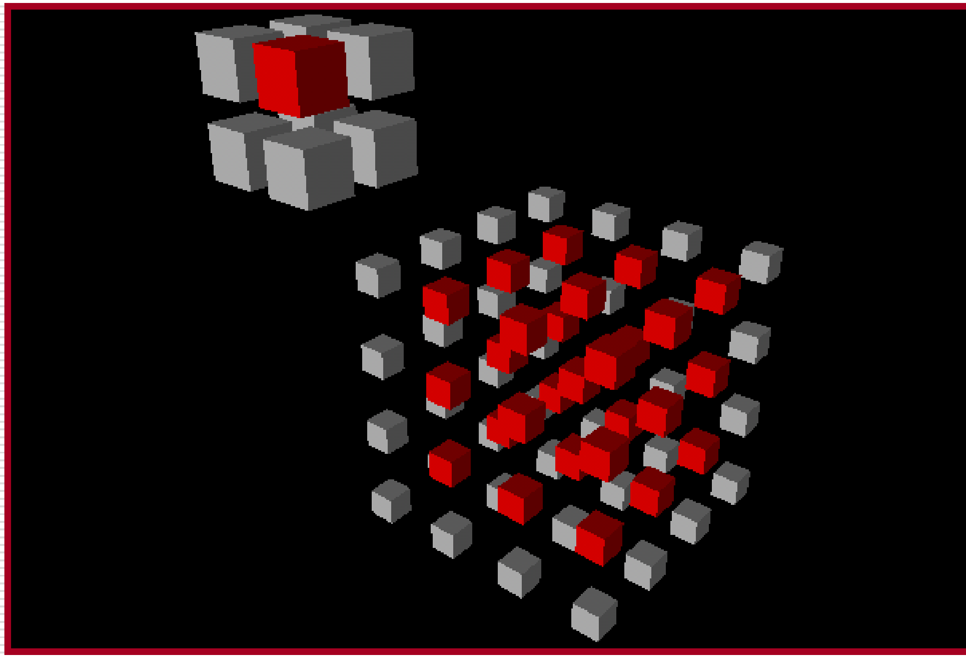

3D-convolution of a 4x4x4 feature map (right) with a 3x3x3 receptive field.

Classification according to the person´s gender – results

(with J. Pihera and J. Veleminska)

Transformation	Error	Standard deviation
2D Raw	0.85%	0.48
2D SOM	14.15%	1.43
3D Direct	8.15%	1.63
3D Direct, clustered	5.37%	1.52
3D SOM	1.28%	0.47

- **Classification according to person´s gender is relatively precise**
- **Raw transformation yields better results for 2D**
- **Pre-processing by a SOM is better for 3D**

Classification of 3D-face models:

accuracy of CNNs around 98% against 64% in humans

(with J. Pihera and J. Veleminska)

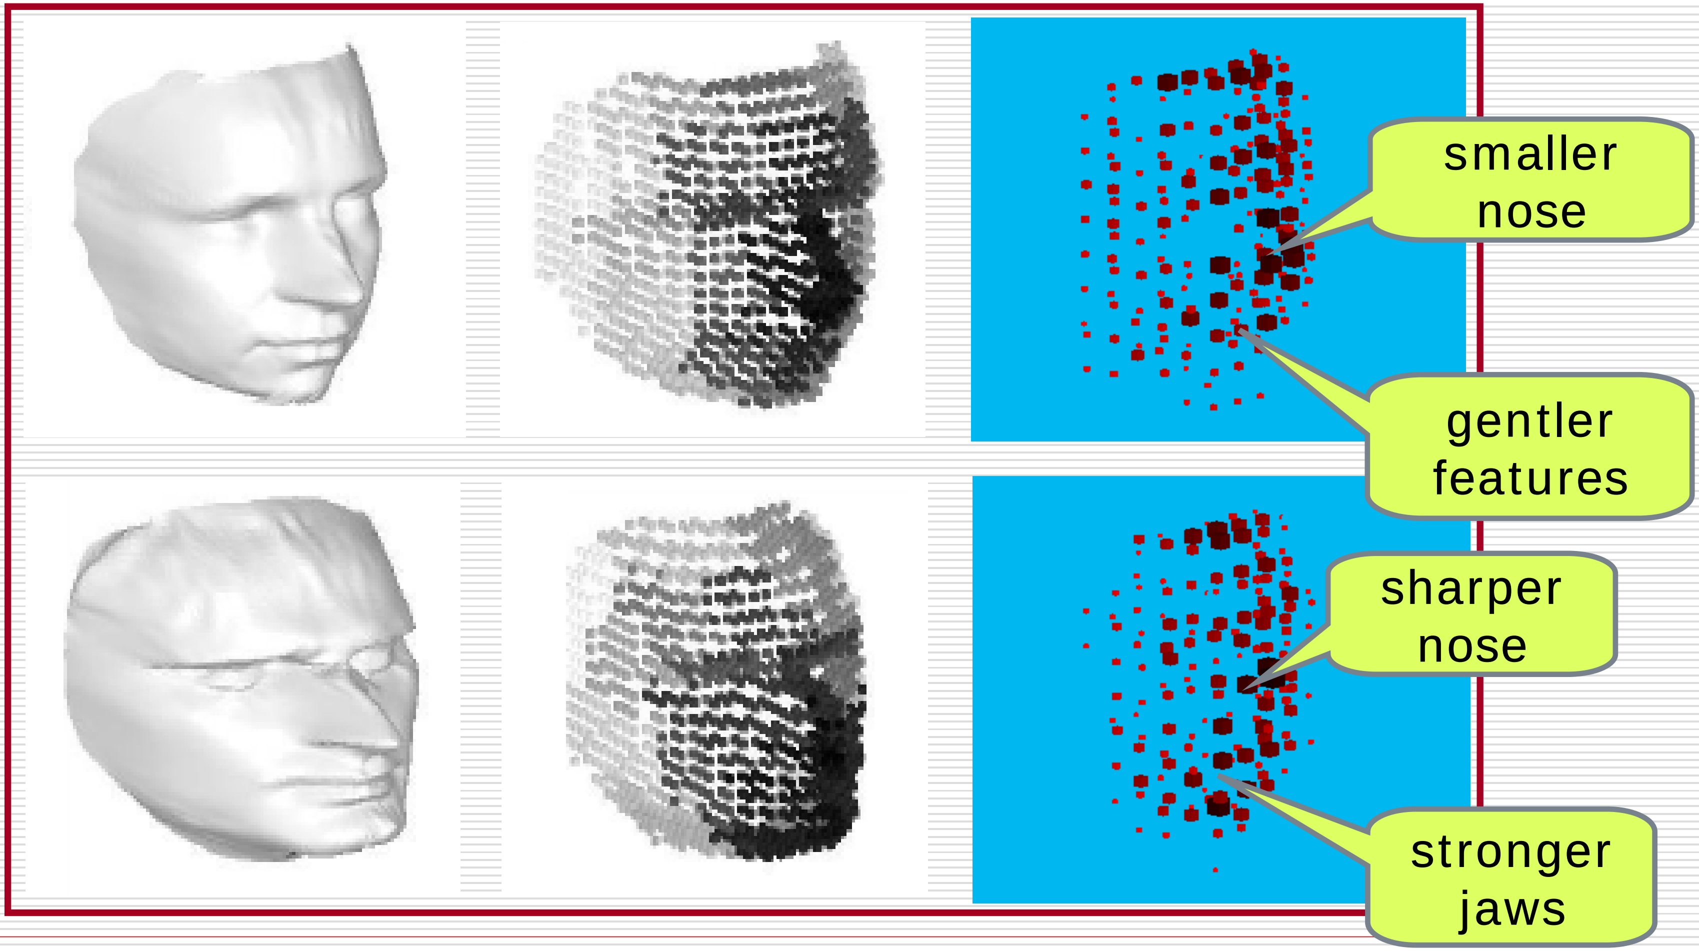

Conclusions

- Understand the function of the brain
- Stimuli for science and industry
- Improved machine performance for at least some tasks should be very welcome

… but shall we really let the machines copy everything from us - even courage, joy, curiosity, …?

Thank you for your attention!

Images courtesy of A. J. Frazer

Acknowledgments: This research was partially supported by the Czech Science Foundation under Grant No. 15-04960S.