Faculty of Mathematics, Physics and Informatics
Comenius University in Bratislava

Neural Networks

Lecture 1

Introduction

Igor Farkaš

2021

Introduction to connectionism

Connectionism – theory of information processing, inspired by biology (the brain). It is based on Artificial Neural Networks (ANNs).

It has two goals:

- theoretical foundations of cognitive science (modeling of cognitive processes)
  - contrasting with symbolic approaches
  - features: parallelism, robustness, learning from experience,...
- applications in practical problems
  - tasks: pattern recognition, classification, associative memory, time series prediction, dimensionality reduction, data visualization, ...

What are the humans good at?

Object recognition

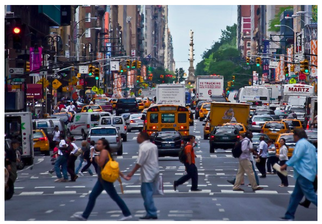

Speech recognition

Natural language

Sensory-motor coordination

Analogical reasoning

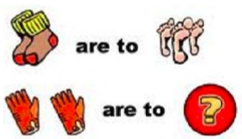

Meanings in general

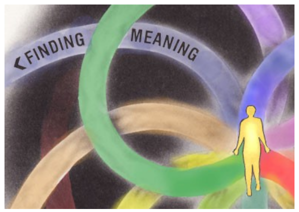

A few facts about human brain

- Brain = highly complex, non-linear and parallel information processing system ("computer metaphor")
- composed of ~ 10^{11} neurons, i.e. brain cells (information-processing elements), connected via ~ 10^{15} synapses
- Glial cells – involved not only in maintenance, but also in information processing
- on certain tasks, brain is much faster than supercomputers of today, even though neurons are very slow (~ ms)
- mostly prewired at birth, but very plastic throughout life
- importance of learning: involves 3 mechanisms
  - modification of existing synapses,
  - generation of new synapses, of new neural cells

Structural organization of levels in the brain

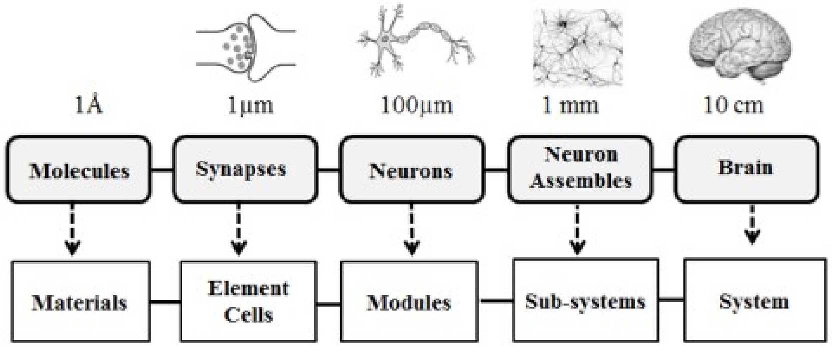

- What is the basic computational level in the brain?
- Not clear, to the contrary with the classical computer

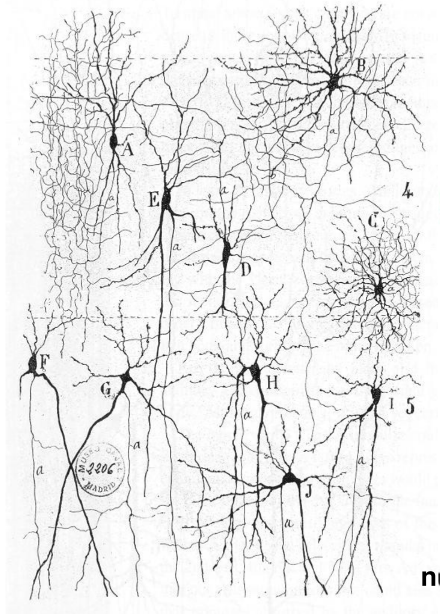

From R. y Cajal: Texture of the Nervous System of Man and the Vertebrates (illustrates the diversity of neuronal morphologies in the auditory cortex).

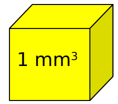

10^5 neurons
3 km axons

Structure of a Typical Neuron

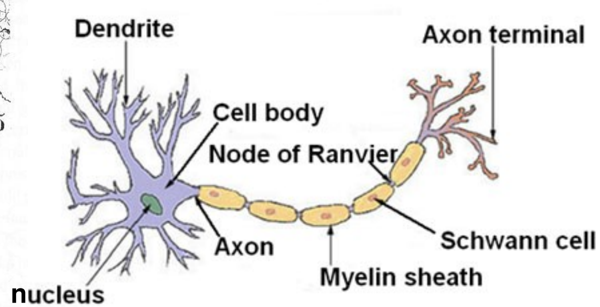

Synapse

- Synapse maintains the interaction between neurons.
- Presynaptic neuron releases a neurotransmitter, which diffuses across the synaptic cleft b/w neurons and then acts on a postsynaptic neuron.
- Synapse mediates electrical-chemical-electrical signal conversion.
- Effect on a postsynaptic neuron can be either excitatory or inhibitory.

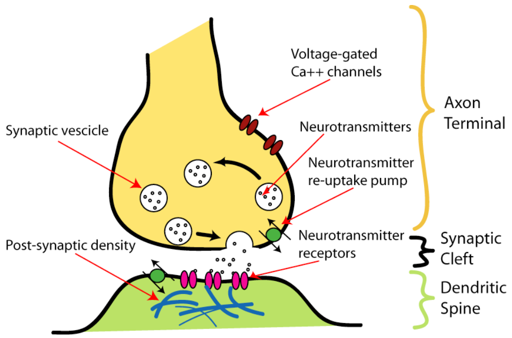

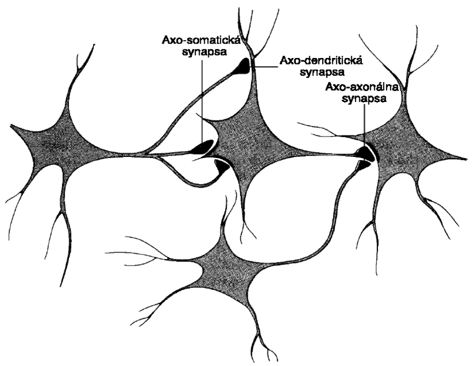

Action potential

If a neuron is made to "fire", generated action potential (AP) traverses along the axon, uninhibited.

Generation of AP (neuron "firing") requires that membrane potential exceed the excitation threshold.

After a spike a neuron recovers (refractoriness).

Each neurons sends out spikes whose frequency can vary.

There exist different theories of neural coding (how do spikes carry information?)

↓

Spiking network models

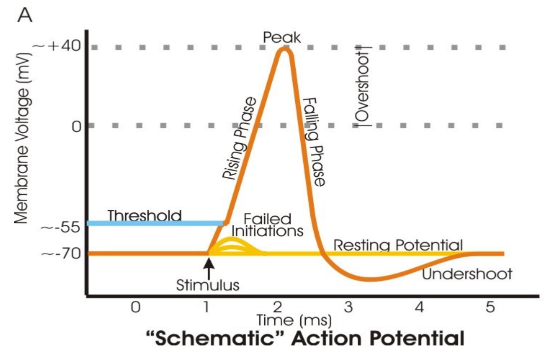

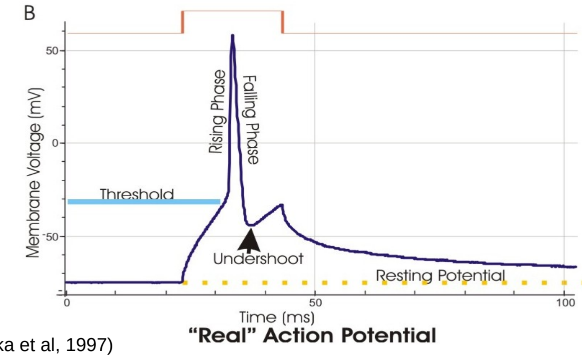

(Kvasnička et al, 1997)

An artificial neuron model

1. receives signals from other neurons (or sensors)
2. processes (integrates) incoming signals
3. sends the processed signal to other neurons (or muscles)

Deterministic model

$o = f(\sum_i w_i x_i - \theta)$

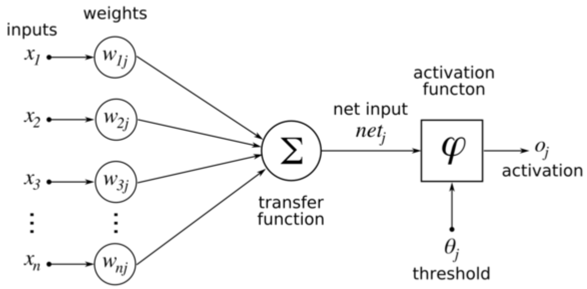

Stochastic model $P(o=1) = 1/(1+\exp(-net/T))$

Features of artificial neural networks

- Nonlinearity (of processing units)
- Input-output mapping (nonparametric statistical inference)
- Adaptivity (parameter tuning)
- Evidential response (degree of 'confidence', soft assignment)
- Contextual information (← thank to connectivity)
- Fault tolerance (graceful degradation)
- VLSI implementability
- Neurobiological analogy
- Uniformity of analysis and design
- Importance of environment (for design)
- New: lack of robustness (against adversarial attacks) ☹

Neural network architectures

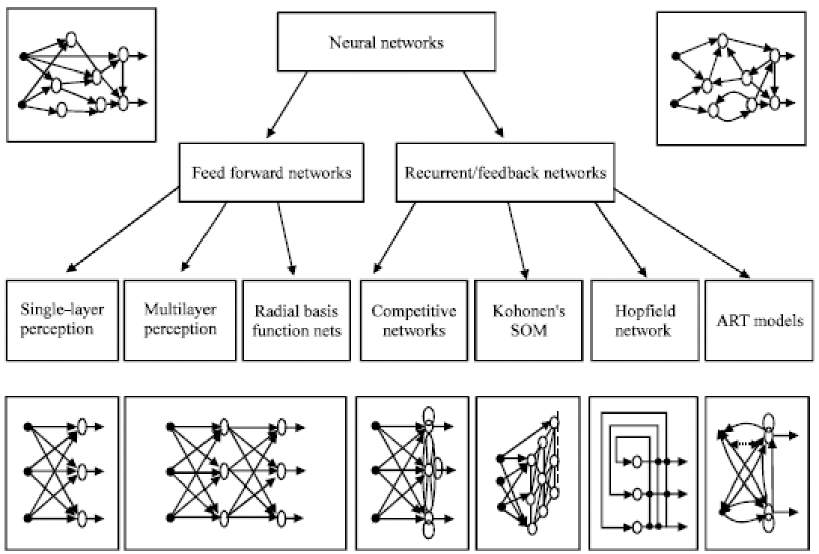

doi:10.3923/itj.2007.526.533

History of classical connectionism

- Aristoteles (400 BC) – introduced concepts of memory, and connectionism
- Spencer (1855) – separated psychology from philosophy, postulated that “neural states affect psychological states”, knowledge is in connections.
- James (1890) – model of associative memory, “the law of habit”
- Thorndike (1932) – distinguished sub-symbolic view on neural associations, formulated two laws of adaptation: “the law of effect” and “the law of exercise” (currently known as reinforcement in operant conditioning).
- McCulloch & Pitts (1943) – neural networks with threshold units
- Minsky (1967) extended their results to comprehensible form, and put them in the context of (formal) automata theory and theory of computation.

First neural network model

- Birth of computer era
- How could information be represented in a nervous system?
- McCulloch & Pitts (1943) – neurons with threshold logic

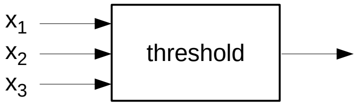

y = 1, if threshold exceeded, otherwise 0

- Weights = 1 (i.e. equal importance of inputs), no learning
- Inhibitory inputs possible (e.g. $y = x_1$ and $\neg x_2$ with threshold 1).
- A single TL unit can simulate any linear Boolean function (BF)
- A two-layer NN with TL units can simulate any BF: $\{0,1\}^n \rightarrow \{0,1\}$
- Birth of neural networks and artificial intelligence disciplines

Milestones of neural networks history

- classical connectionism (until 1940s)
  - within philosophy, psychology
- 1st NN wave (1940-1960) – birth of computer era, cybernetics
  - beginning of theory of ANN, linked to cognitive science revolution
- 2nd NN wave (1980-1995)
  - parallel distributed processing → subsymbolic processing
  - multi-layer NN models (incl. recurrent)
  - Later: multilayer generative models (probabilistic approach)
- 3rd NN wave - renaissance of ANNs (2006-)
  - deep networks, convolutional NN, reservoir computing

Knowledge representation

- Knowledge refers to stored information or models used by a person or machine to interpret, predict and appropriately respond to the outside world. (Fischler & Firschein, 1987)
- Goal of NN learning: learn the task (model) and maintain it.
- training examples – labeled or unlabeled
- KR is goal oriented: In “intelligent” machines, a good solution depends on a good KR.
- ANNs are a special class of intelligent machines.
- (Long-term) knowledge in ANN is distributed in free parameters (synaptic weights)
  - but NN architecture also enables knowledge creation

Roles of knowledge representation

1. Similar inputs (patterns) should usually produce similar representations inside a NN, and should hence be classified as the same category (w.r.t. to a similarity measure: e.g. Euclidean distance).
2. Items to be categorized as separate classes should be given widely different representations in a NN.
3. Important features should involve a larger number of neurons for its representation.
4. Prior information and invariance should be built into the design of a NN, hence simplifying the learning task:
   - Prior: restricting NN architecture, constraining the weights
   - Invariance: by structure, by training, by feature space

Representation learning – in context of deep learning.

Subfields of AI

- DL – towards end-to-end learning
- leads to representation learning
- Computational intelligence (see wiki)
- also subset of AI, focusing on soft computing

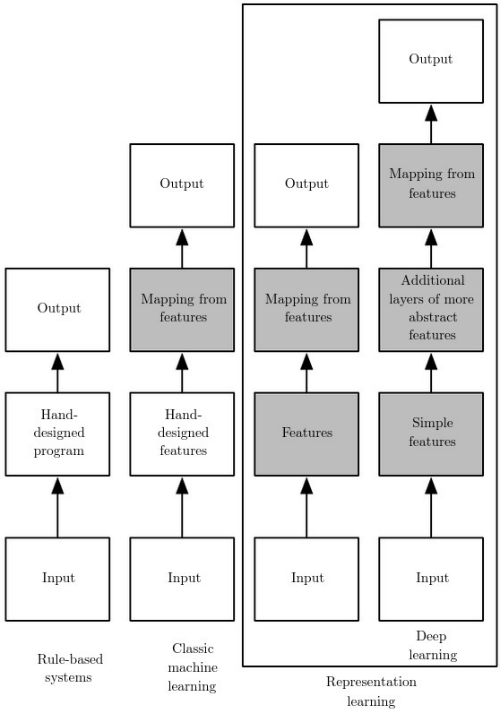

(Goodfellow et al, 2015)

Significant recent developments

- Deep learning in feedforward multilayer networks – very successful in various domains (image recognition, speech processing, language modeling)
- Reservoir computing – efficient approach to processing spatio-temporal signals
- NNs as building blocks of various cognitive architectures
- Computational neuroscience (spiking neural networks) – quest for neural code
  - Blue Brain Project (since 2005)
  - (huge EU) Human Brain Project (since 2013)
- neuromorphic computing
- Trend towards explainable (i.e. trustworthy) AI

Learning paradigms in NN

supervised (with teacher)

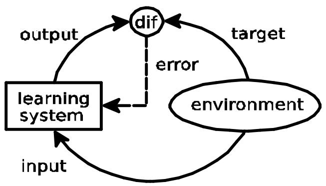

unsupervised (self-organized)

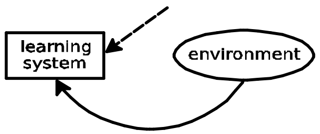

reinforcement learning (partial feedback)

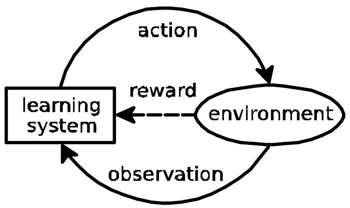

Learning rule types in ANN

- Error-correction – supervised
  - closed-loop feedback system
- Memory-based (e.g. k-nearest neighbors classifier)
  - knowledge stored in examples
- Hebbian – unsupervised
  - correlational synaptic potentiation/depression
- Competitive – unsupervised
  - competition for inputs, feature detectors
- Boltzmann – stochastic
  - inspired by from stat. mechanics, good for high-dim. problems

Learning tasks

- Pattern association (auto-, hetero-)
- Pattern classification (within pattern recognition)
- Feature extraction (within PR or independently)
- Data compression
- Data visualization
- Function approximation
- Control
- Filtering
- Prediction
- Signal generation (with recurrent networks)
- ...

Progress in HW over decades

Decade	Dataset	Memory	Floating Point Calc / sec.
1970	100 (Iris)	1 KB	100 KF (Intel 8080)
1980	1 K (House prices in Boston)	100 KB	1 MF (Intel 80186)
1990	10 K (optical character recognition)	10 MB	10 MF (Intel 80486)
2000	10 M (web pages)	100 MB	1 GF (Intel Core)
2010	10 G (advertising)	1 GB	1 TF (Nvidia C2050)
2020	1 T (social network)	100 GB	1 PF (Nvidia DGX-2)

(Zhang et al, 2020)

Model building

- Key components: data, model, algorithm
- In supervised learning: objective (cost, loss) function

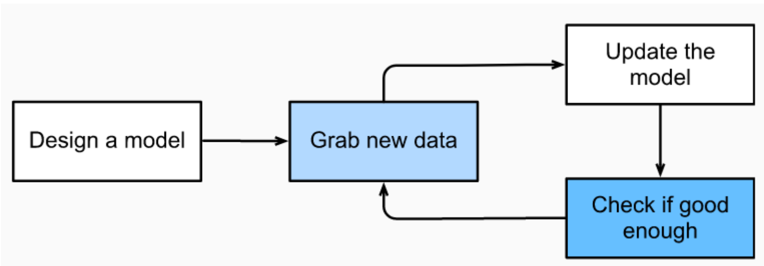

(Zhang et al, 2020)

Conclusion

- Artificial neural networks are a brain inspired computational approach toward solving various complex tasks.
- Each task can be viewed as a (typically nonlinear) mapping between input and output representations.
- Representations have be in a numeric format.
- During the last decade, we have witnessed the third wave of successful ANN applications, mainly due to deep learning, within the scope of narrow AI.
- There are still many challenges for ANN to reach human-level intelligence.
- Prediction: Bio-inspired modeling may be crucial.