Nizhny Novgorod State University

Institute of Information Technologies, Mathematics and Mechanics

Department of Computer software and supercomputer technologies

Educational course
«Introduction to deep learning
using the Intel® neon™ Framework»

Introduction to deep learning

Supported by Intel

Valentina Kustikova,
Phd, lecturer, department of Computer software
and supercomputer technologies

Content

- What is “deep learning”?
- History
- Practical problems effectively solved using deep learning
- Biological neurons and its similarity with artificial neurons
- Model of artificial neuron
- Deep model classification based on supervised and unsupervised principle

What is “deep learning”?

- ***Deep Learning*** is a field of machine learning, that examines methods for solving problems of artificial intelligence using deep neural networks

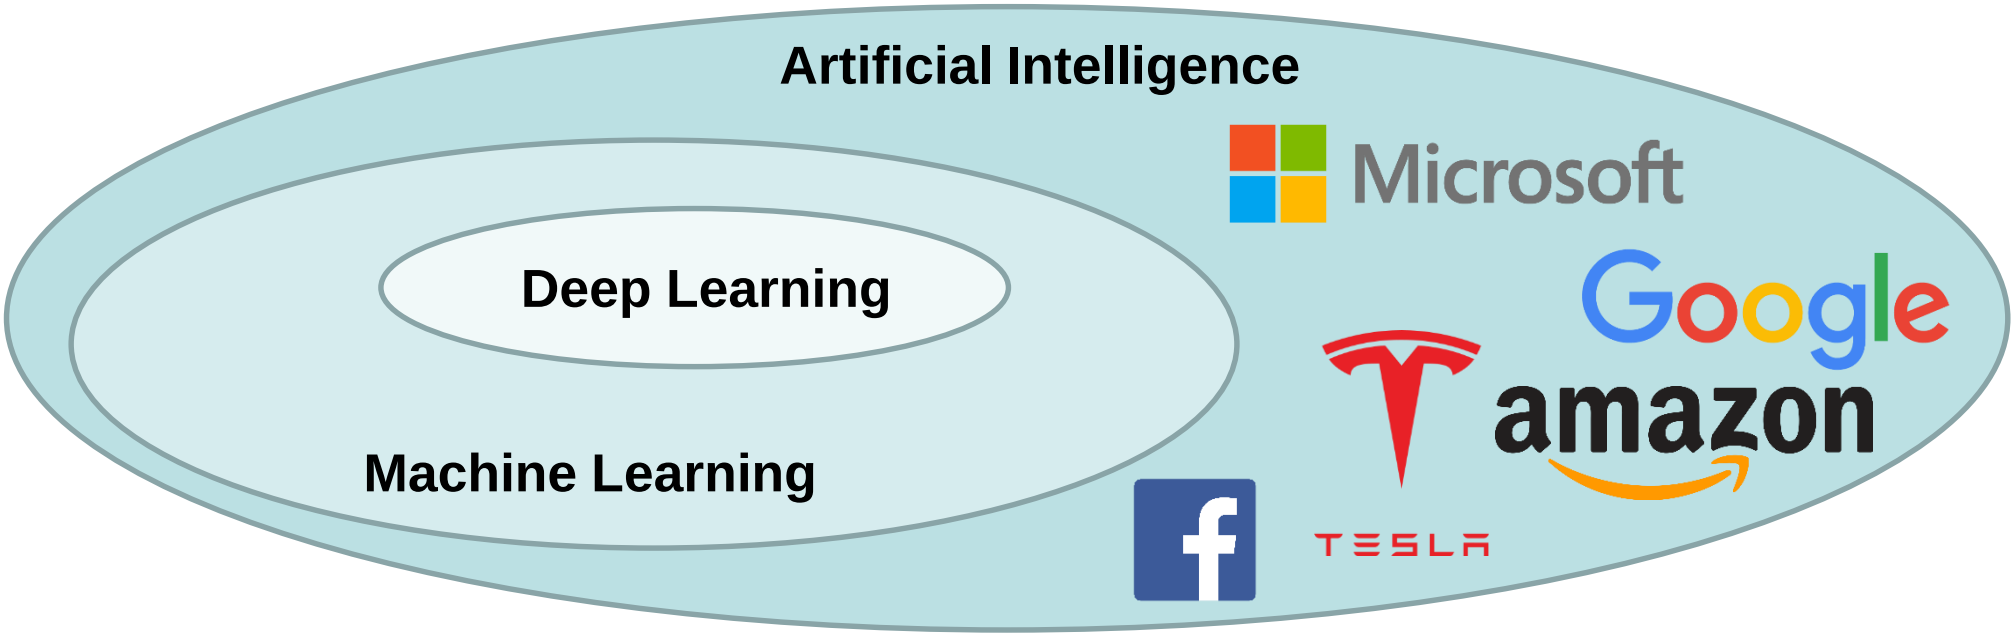

History

1. The complexity of training deep models
2. Lack of necessary computational resources
3. Small amounts of training datasets

Effective implementation of training deep neural networks

1957	1960x	1986	**2006**	2011	2012
Rosenblatt's Perceptron	Mass selection method	Back Propagation Method	Development of GPUs and multicore systems	Deep Belief Networks	Deep Convolutional Neural Networks
				The problem of speech recognition	*The problem of image classification (1000 categories)*
				Microsoft Research	G. Hinton with colleagues


F. Rosenblatt

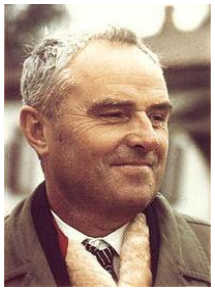
A.G. Ivakhnenko


G. Hinton

Practical problems (1)

- The most famous samples of successful practical application of deep learning:
  - The artificial intelligence AlphaGo defeated the best player in Go

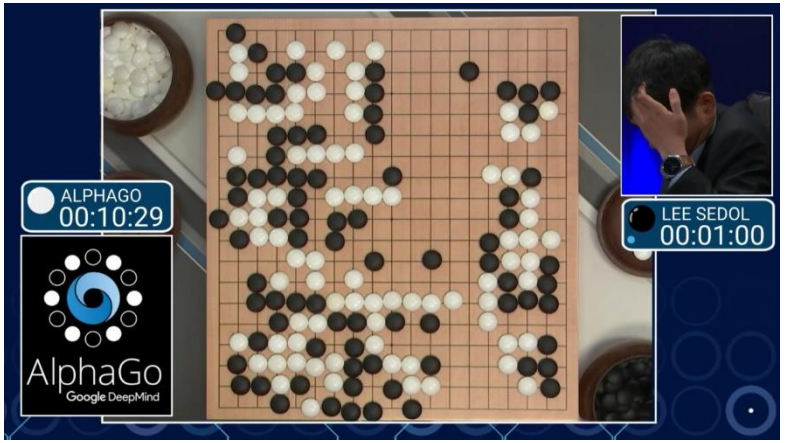

  - Technologies of autonomous driving cars (Google, Tesla, Uber)

Practical problems (2)

- Other examples of practical application of deep training:
  - Recommended system for users of the online store Amazon
  - Recommended system for users of the Netflix service
  - Google voice search
  - Personal assistants Alexa developed by Amazon and Cortana developed by Microsoft. Personal assistant accepts voice commands to create to-do lists, order items online, set reminders, etc.
  - Deep face recognition (DeepFace) developed by Facebook
  - …

Practical problems (3)

- The following examples are considered:
  - Problems in natural language processing (online translators, text generators)
  - Problems in computer vision (image classification, object detection, semantic segmentation)

Online translators

- Google Translate

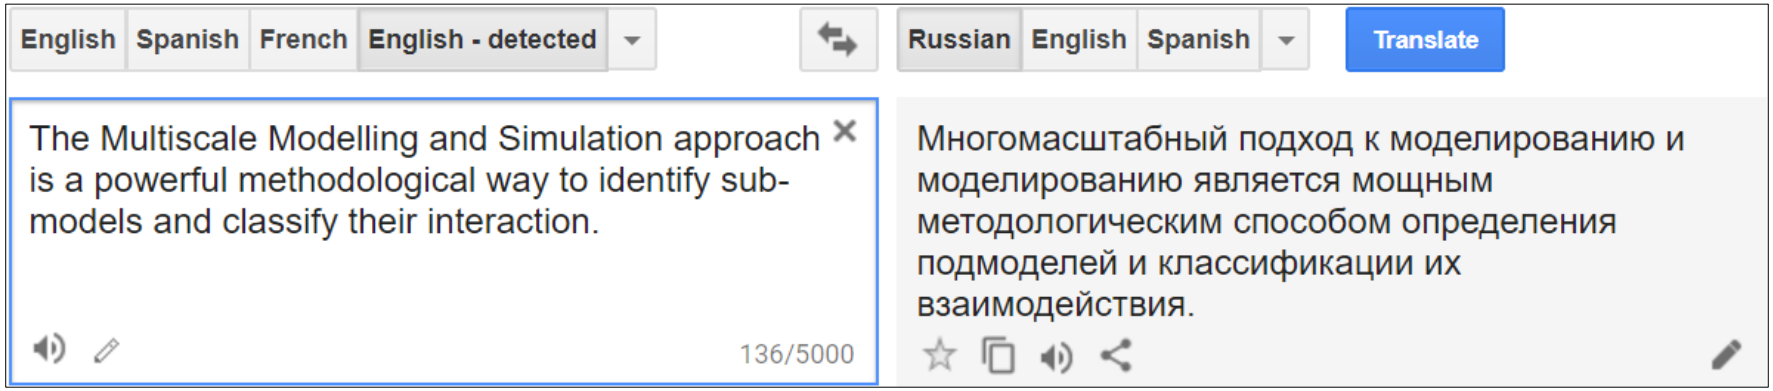

- Yandex Translator (Яндекс.Переводчик)

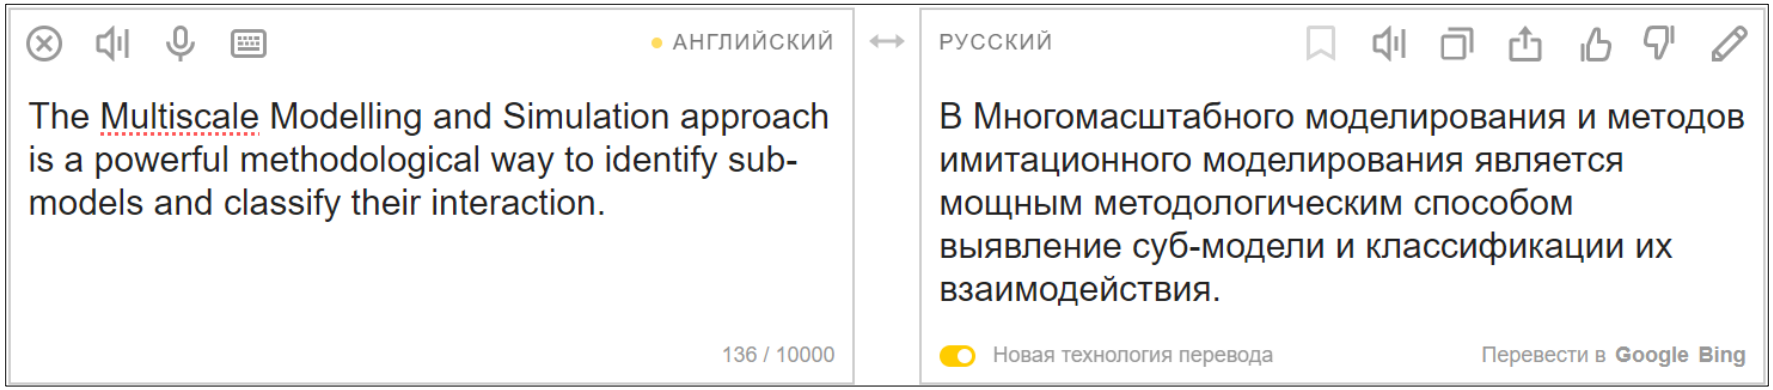

Text generators

- ***Text generators*** are programs that provide automatic text generation, correct from the point of view of most language norms, but, as a rule, meaningless
- Text generators are used in the development of virtual interlocutors (chatbots and commentators in social networks and blogs)

- 15 Dummy Text Generators You Should Know [https://www.webdesignerdepot.com/2012/03/15-dummy-text-generators-you-should-know]

Image classification problem

- The problem of image classification is to match image to the represented object class

* Russakovsky O., Deng J., Su H., Krause J., Satheesh S., Ma S., Huang Z., Karpathy A., Khosla A., Bernstein M., Berg A.C., Fei-Fei L. ImageNet Large Scale Visual Recognition Challenge // International Journal of Computer Vision, 2015.

Image classification on ImageNET dataset

Year	Team	Method	Misclassification error*
2010	NEC-UIUC (Tokyo, Japan)	Descriptor Coding + SVM	0.28191
2011	XRCE (Xerox Research Center Europe, Cordoba University, Argentina)	Fisher Vectors + one-vs-all linear SVMs	0.25770
2012	SuperVision (University of Toronto, Canada)	Convolutional Neural Network (AlexNet)	**0.15315**
2013	Clarifai	Multiple Neural Networks	0.11197
2014	GoogLeNet	Convolutional Neural Network (GoogLeNet)	**0.06656**
2015	MSRA	Deep Residual Network	0.03567
2016	Trimps-Soushen (The Third Research Institute of the Ministry of Public Security, P.R. China)	Ensemble of Convolutional Neural Networks	0.02991

Note: misclassification error is a relation of correct classified images to the overall number of test images.

Object detection problem

- The problem of object detection is to determine the placement of bounding boxes for the specific object class

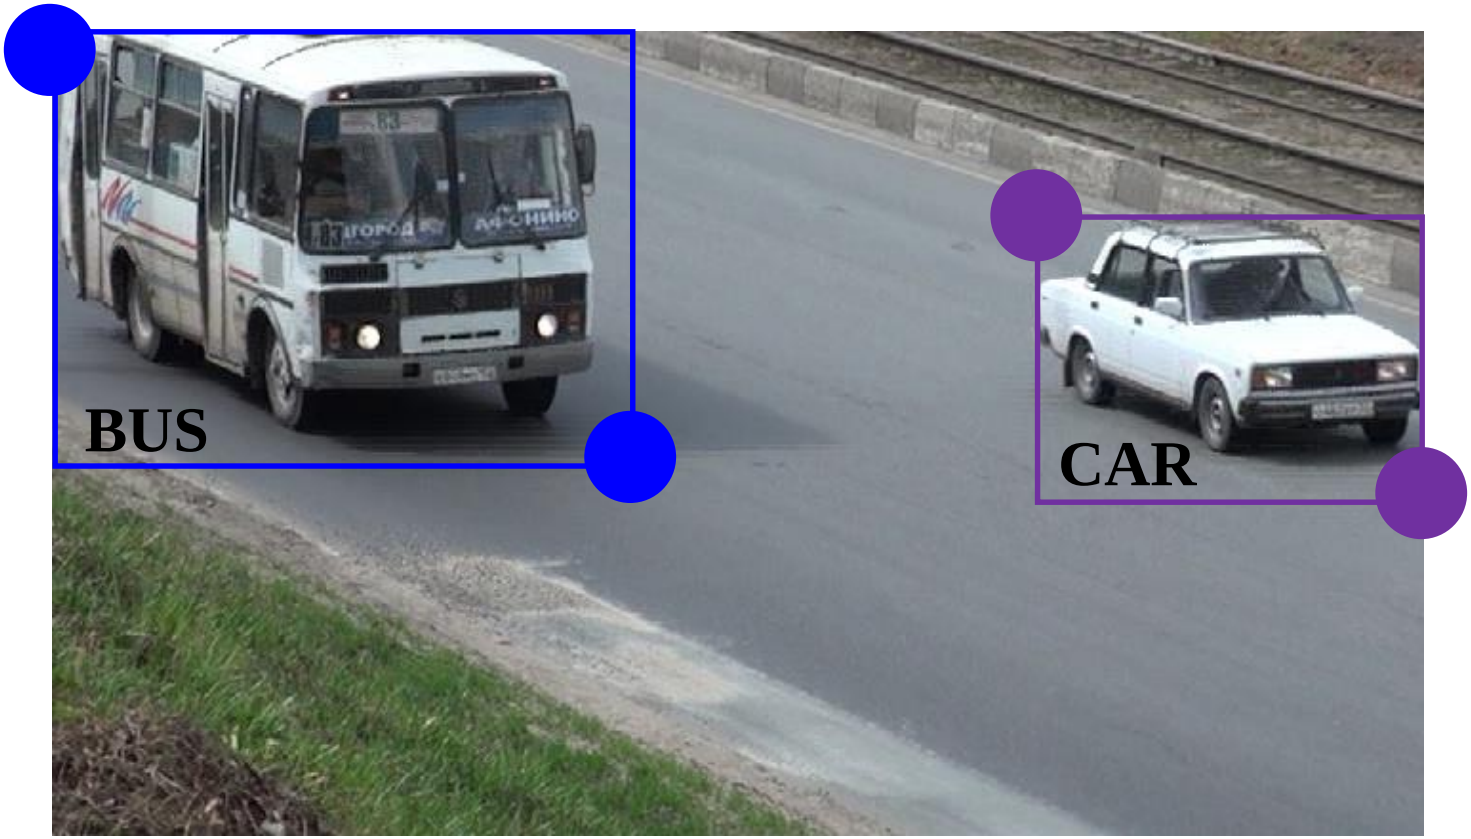

Object detection results (average precision) on PASCAL VOC 2012 dataset

Object class	LSVM[1] is a deformable part-based model (2011)	YOLO[2] is a single neural network (2016)	Difference
BOTTLE	18.1	18.8	0.7
CAT	24.2	65.6	**41.4**
DINING TABLE	4.5	35.9	31.4
DOG	17.5	61.4	**43.9**
HORSE	15.2	57.9	**42.7**
PERSON	7.9	63.8	**55.9**
SOFA	7.1	39.5	32.4
TV/MONITOR	25.7	46.2	20.5
Average on 20 classes	**20.9**	**48.8**	**27.9**

Note: the average precision reflects the correctness of the constructing bounding boxes, so the higher the value, the better the method detects objects.

1. Felzenszwalb P.F., Girshick R.B., McAllester D., Ramanan D. Object Detection with Discriminatively Trained Part Based Models // IEEE Transactions on Pattern Analysis and Machine Intelligence. – 2010. – Vol. 32, No. 9.
2. YOLO: Real-Time Object Detection [https://pjreddie.com/darknet/yolo].

Semantic segmentation problem

- The problem of semantic segmentation is to map each pixel of an image to the object class to which it belongs

Origin image | Layout | Segmentation result

* The PASCAL Visual Object Classes Homepage [http://host.robots.ox.ac.uk/pascal/VOC].

Semantic segmentation results (average precision) on PASCAL VOC 2012 dataset

Object class	DeepLab-CRF (Deep Convolutional Neural Network + Conditional Random Fields) (2014)	SYSU_SceneParsing_COCO, ResNet-101 (2016)	Difference
AEROPLANE	78.4	94.6	16.2
BICYCLE	33.1	66.7	**33.6**
CHAIR	25.3	52.3	27
COW	69.2	94.9	25.7
DINING TABLE	52.7	75.8	23.1
DOG	75.2	93.2	18
HORSE	69.0	95.5	26.5
SOFA	45.1	78.4	**33.3**
TV/MONITOR	56.2	94.6	**38.4**
Average on 20 classes	**66.4**	**85.7**	**19.3**

Note: average precision of the semantic segmentation reflects the number of correctly classified pixels. In this case, the pixels belonging to the object boundaries are not taken into account.

Style transfer

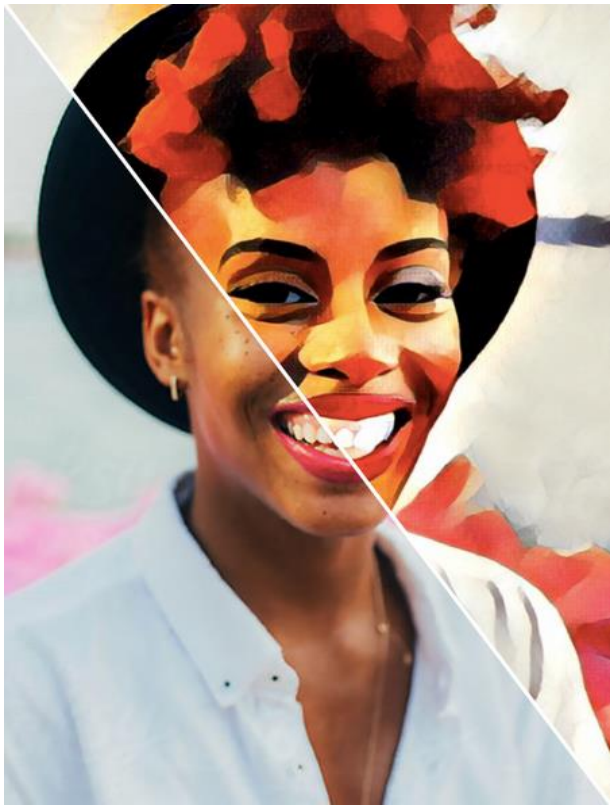

* PRIZMA Labs. Superior Image Analysis [https://prismalabs.ai/api-sdk.html#style-transfer].

Colorizing B&W photos and videos with neural networks

* 30 amazing applications of deep learning [http://www.yaronhadad.com/deep-learning-most-amazing-applications].

Biological neurons and its similarity with artificial neurons (1)

- Artificial neural network models processing information by a human brain:

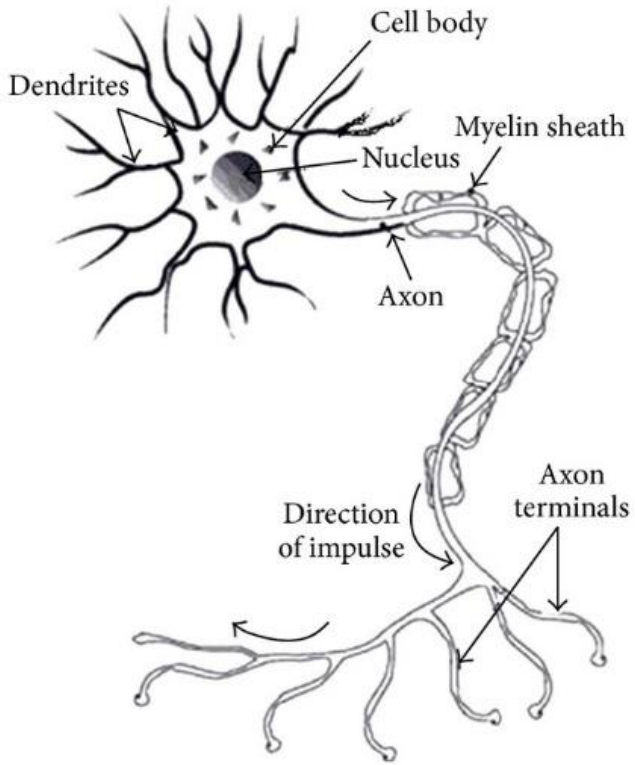

* Ching Lee Koo, et al. A Review for Detecting Gene-Gene Interactions Using Machine Learning Methods in Genetic Epidemiology.

Biological neurons and its similarity with artificial neurons (2)

- The mechanism of the functioning of biological neurons explains the behavior of neurons in an artificial neural network
- The biological neuron consists of the following parts:
  - The body, called ***soma***, where the core is located
  - There are two types of appendages: thin densely branched ***dendrites*** and an ***axon***. Axon is a thicker appendage, splitting at the end
  - Input signals enter the cell through ***synapses***
  - The output signal is removed by the axon through the nerve endings, called the ***collaterals***

The model of artificial neuron

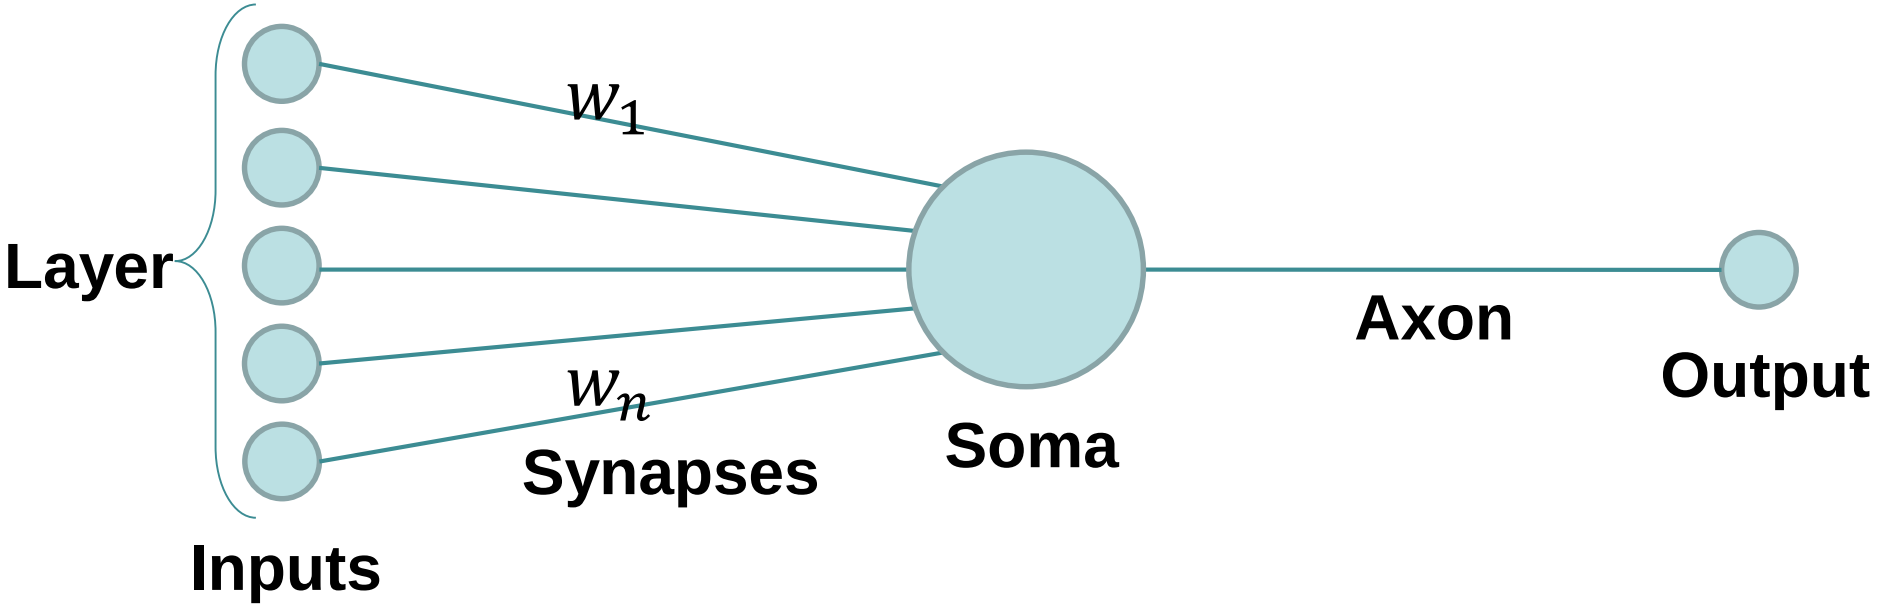

- Axons of input neurons represent synapses of the current neuron, the axon of this neuron is a synapse of the output neuron
- Neurons of the same level form a layer
- Training of the neural network means tuning the weights of synaptic channels

Deep model classification

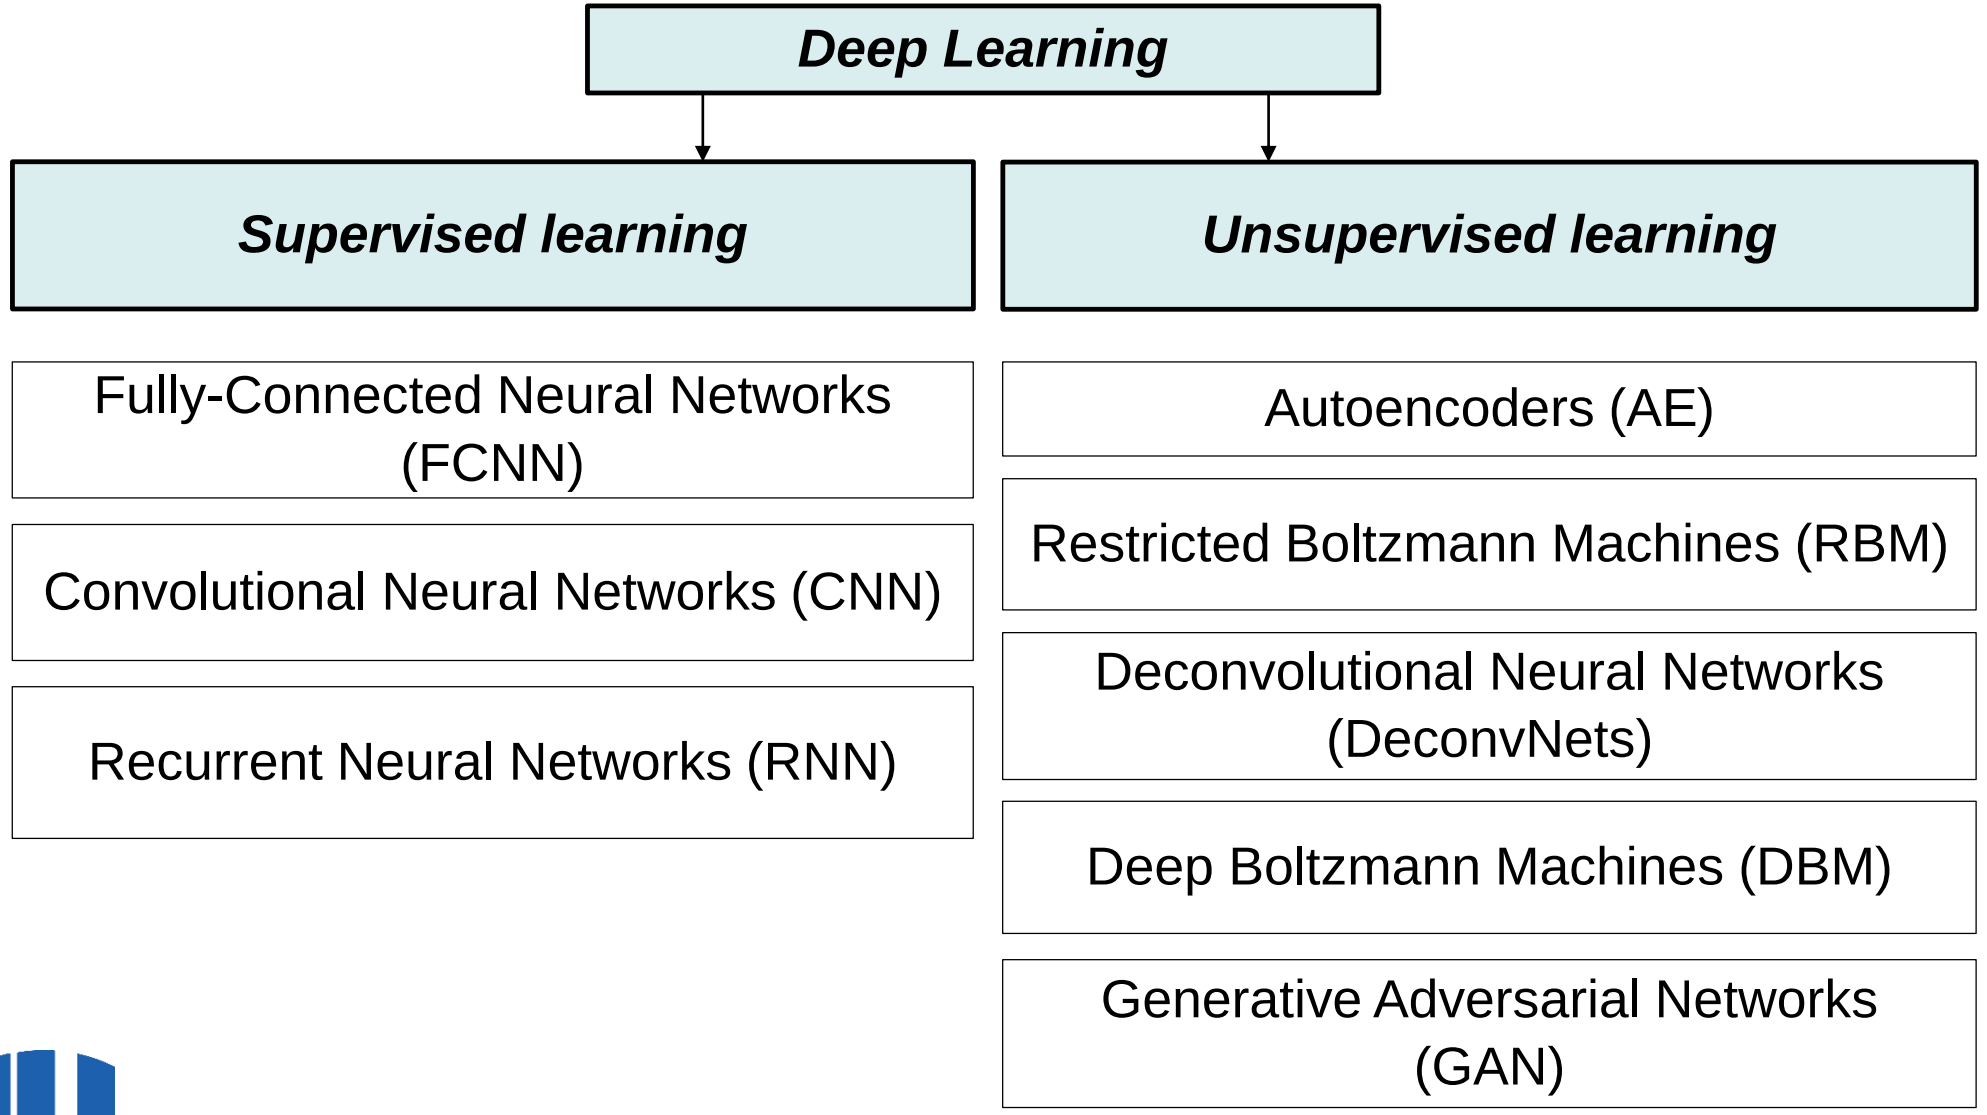

Deep model classification. Models considered in the course in details…

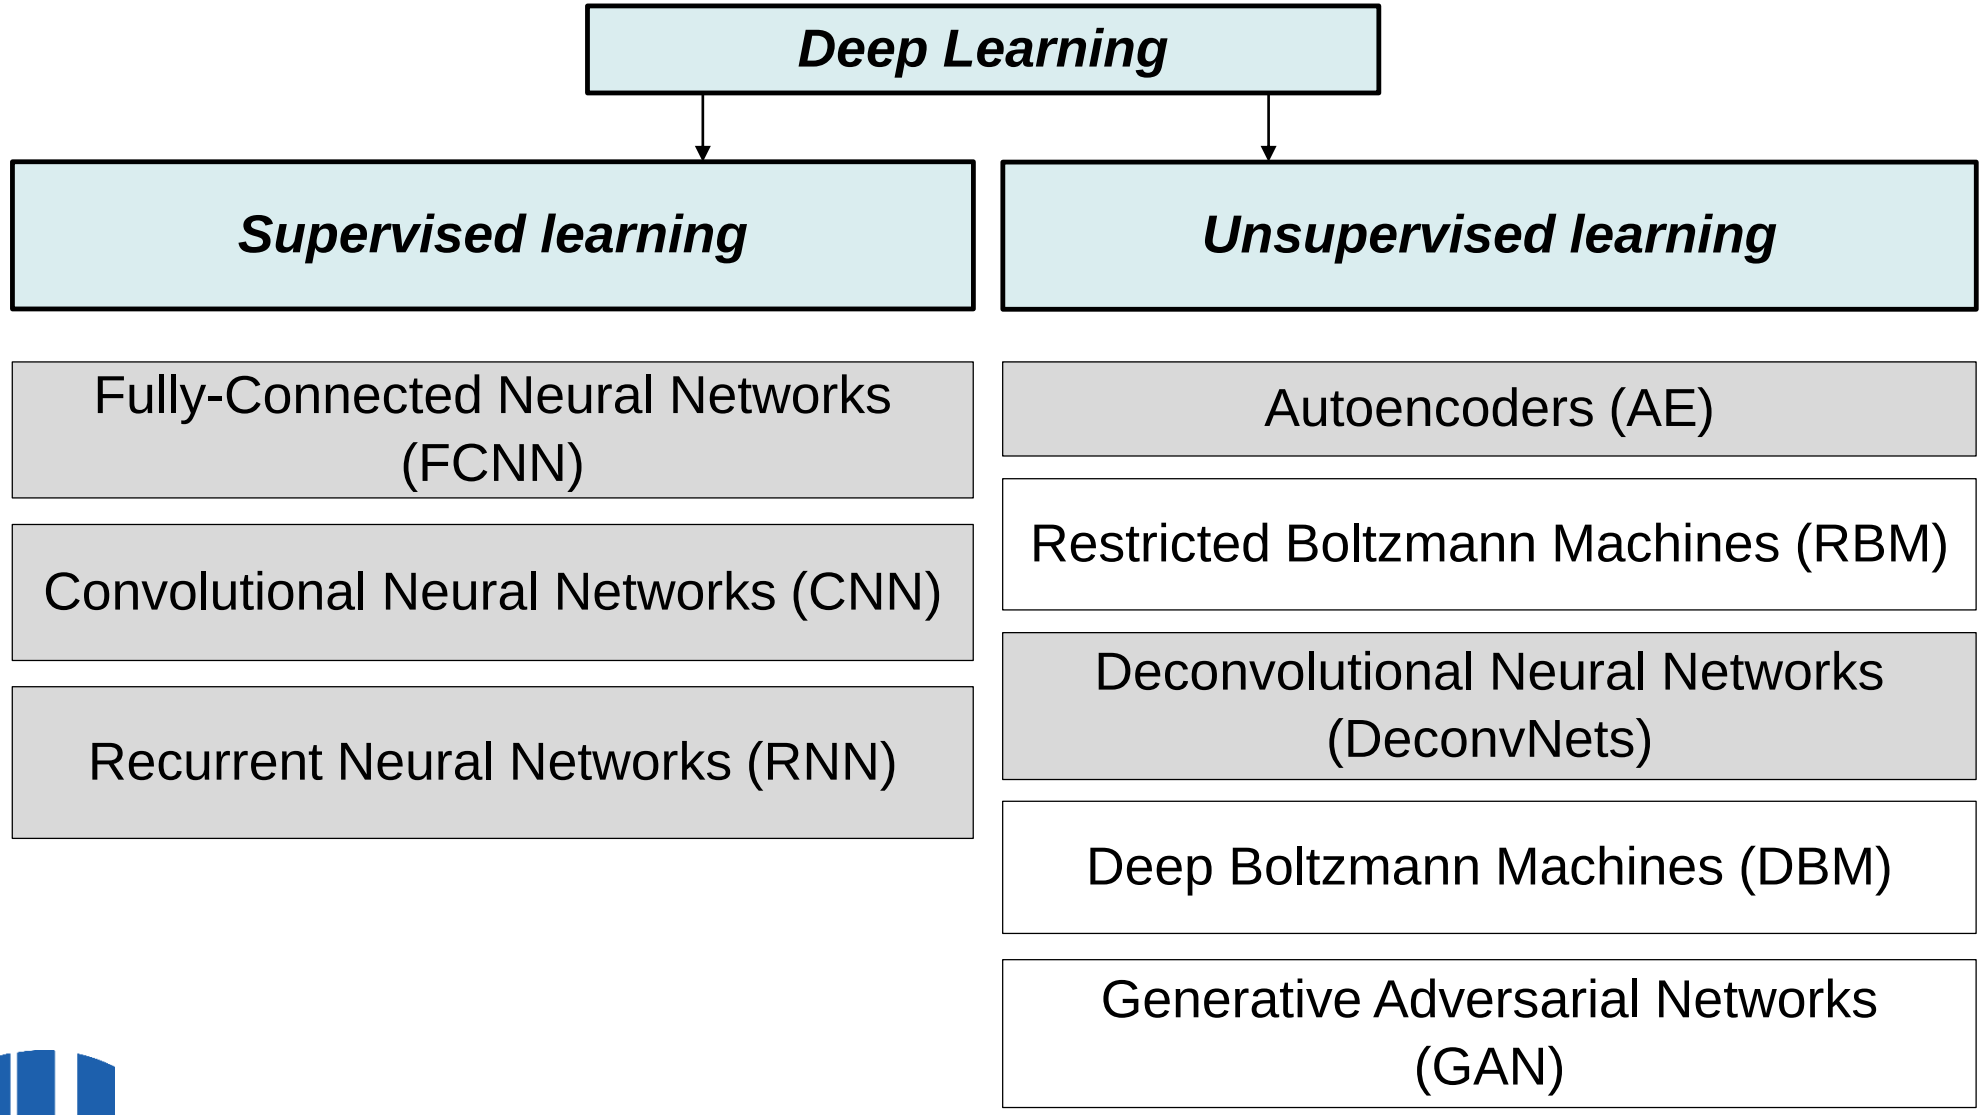

Deep model classification. Models considered in the course…

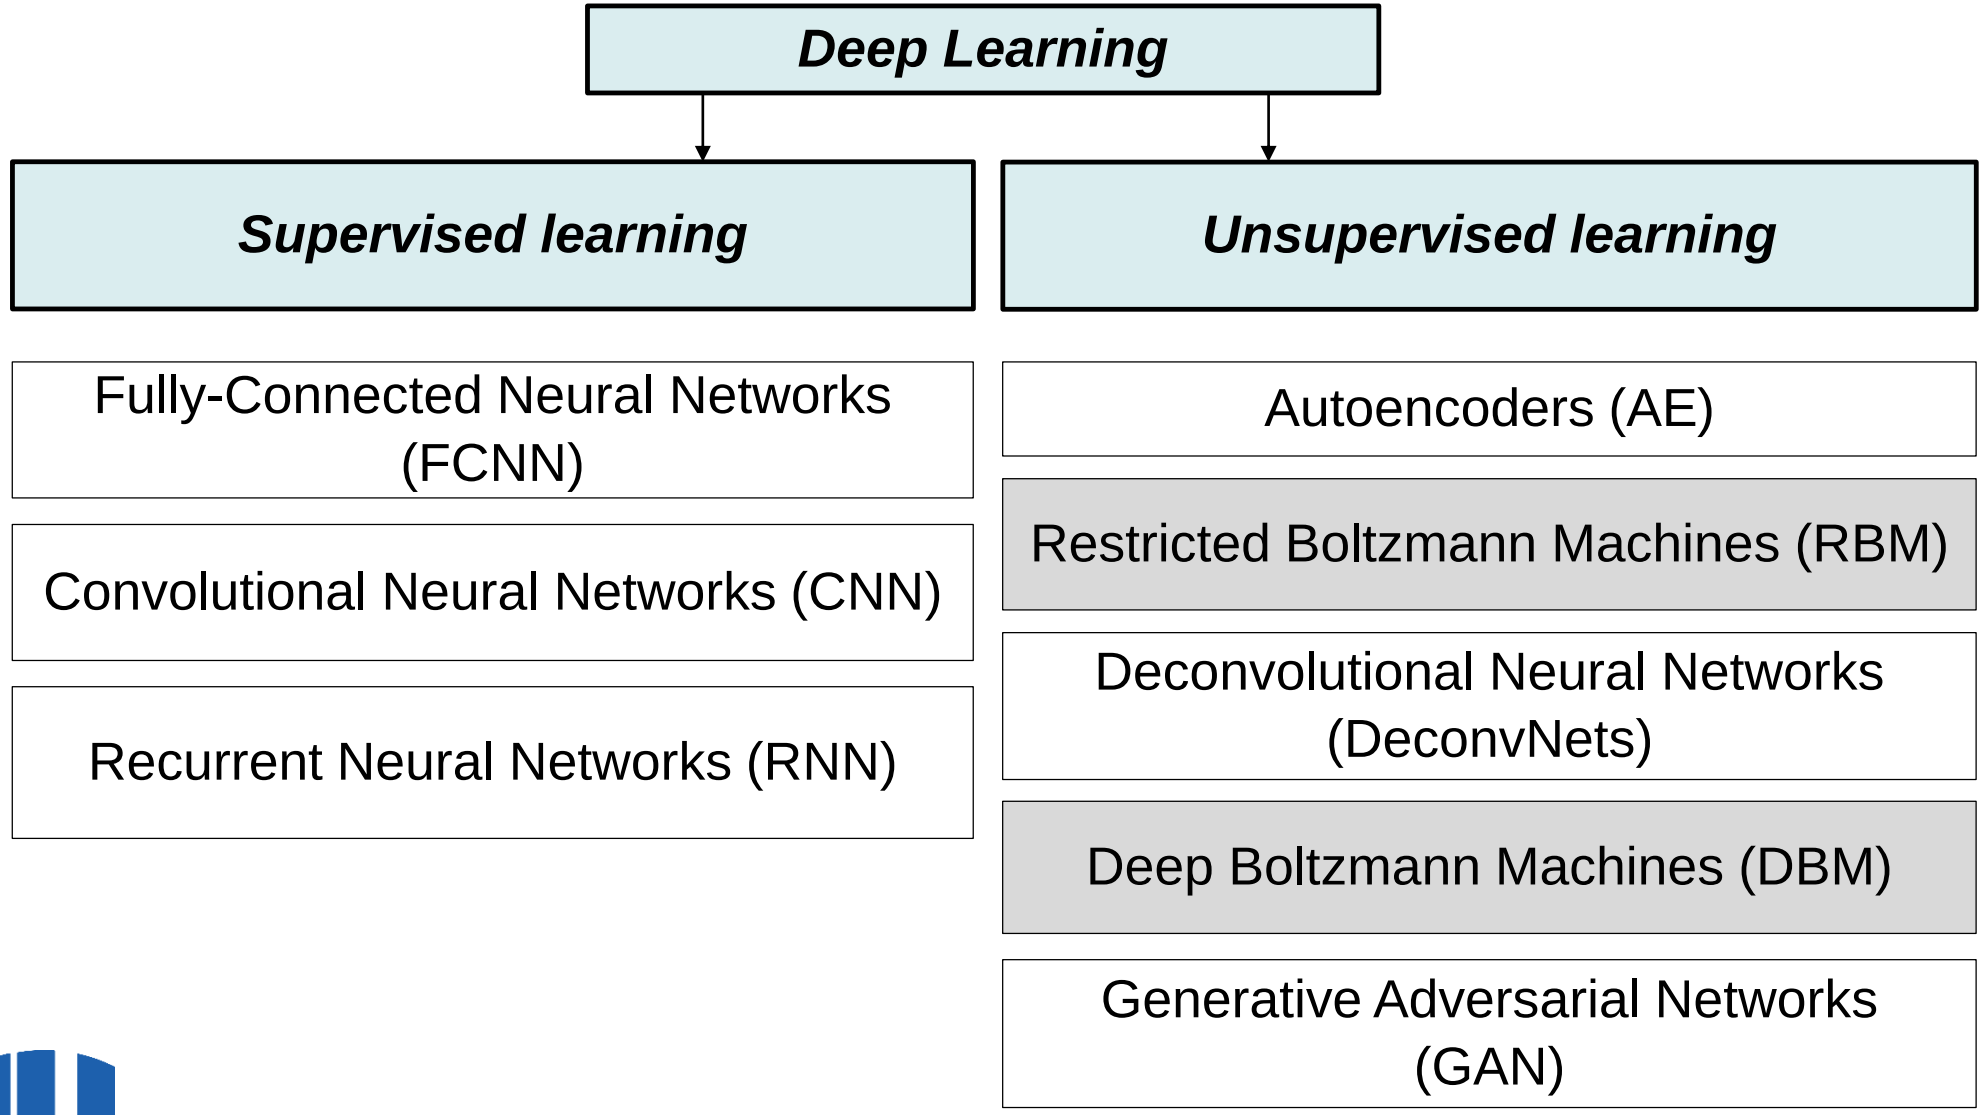

Conclusion

- Deep learning has wide practical use in various areas
- In the course the basic deep models and their application are considered on the example of one task of computer vision
- Implementation of the developed models is performed using Intel® neon™ Framework

Literature

- Haykin S. Neural Networks: A Comprehensive Foundation. – Prentice Hall PTR Upper Saddle River, NJ, USA. – 1998.
- Osovsky S. Neural networks for information processing. – 2002.
- Goodfellow I., Bengio Y., Courville A. Deep Learning. – MIT Press. – 2016. – [http://www.deeplearningbook.org].

Authors

- **Kustikova Valentina Dmitrievna**
  Phd, lecturer, department of Computer software and supercomputer technologies, Institute of Information Technologies, Mathematics and Mechanics, Nizhny Novgorod State University
  valentina.kustikova@itmm.unn.ru
- **Zhiltsov Maxim Sergeevich**
  master of the 1st year training, Institute of Information Technology, Mathematics and Mechanics, Nizhny Novgorod State University
  zhiltsov.max35@gmail.com
- **Zolotykh Nikolai Yurievich**
  Dr., Prof., department of Algebra, geometry and discrete mathematics, Institute of Information Technologies, Mathematics and Mechanics, Nizhny Novgorod State University
  nikolai.zolotykh@itmm.unn.ru