iTMO

ITMO University

Advanced Machine Learning

Course Workload		Assessment form (examination/ graded test/ ungraded test)
ECTS	Hours	
5	180	Exam + course project

The course examines modern algorithms for constructing deep neural networks that allow solving applied problems related to image and video processing, natural language analysis, generation of new objects, as well as reinforcement learning methods. Lectures are accompanied by intensive laboratory exercises devoted to the practical application of the algorithms discussed in the theoretical part of the course.

Course structure:

1. Neural network architectures

1.1. McCulloch-Pitts model. Rosenblatt's perceptron. Neural implementation of logical functions.

1.2. Activation functions and error functions. Error functions for classification and regression problems.

1.3. Optimization problem statement. Gradient and non-gradient optimization methods.

1.4. Gradient descent and its types. Stachastic gradient descent and mini-batch method.

1.5. Gradient descent modifications: the method of moments, the RMSProp algorithm and the Adam algorithm.

1.6. Methods for initializing neural network weights.

1.7. Multilayer neural network. Backpropagation method. The problem of damped and exponentially growing gradients.

1.8. Neural network training as an optimization problem. Stochastic gradient descent with moments, RMSProp and Adam algorithms.

1.9. Training data: train, test, sample validation. Monitoring the learning process of the network.

1.10. Using parallel computing in deep learning.

1.11. Parallel optimization of neural network hyperparameters.

1.12. Convolutional neural network architecture. The principle of "shared weights", folding, pooling, padding.

1.13. Convolutional neural network architecture. The principle of "shared weights", folding, pooling, padding.

1.14. Modern convolutional networks. VGG-16 and ResNet models. Attention mechanism. Problems of convolutional networks and capsule networks.

1.15. Sequence Analysis Problems. The principle of constructing recurrent neural networks.

1.16. Recurrent neural networks. LSTM and GRU cells. Attention mechanism.

2. Applied problems of modern machine learning

2.1. Natural language processing tasks. A bag of words approach. TF-IDF method. Stamming. Lemmatization. Stop words.

2.2. Vector word representation and embeddings. The word2vec model.

2.3. Topic modeling methods. LDA method. Visualization of the results of thematic modeling.

2.4. Applied thematic modeling. Assessment of the quality of highlighting topics in the corpus of documents.

2.5. Autoencoder architecture. Encoder and decoder.

2.6. Autoencoders and code space. Generation of new objects. Variational autoencoders.

2.7. Application of the attention mechanism in neural networks. "Multi-head attention".

2.8. Transformers architecture ("Attention is all you need").

2.9. Generative adversarial networks. GAN learning algorithm.

2.10. Style transfer. "Style error" and "content error".

2.11. Conditional GAN architecture. Adding a condition to the input of the model.

2.12. Capsule neural networks and inverse graphics.

2.13. Algorithms and architectures of neural networks for 3D detection of objects in the image.

2.14. Application of deep learning methods in the task of developing unmanned vehicles.

2.15. GANs architectures for generating 3D objects.

2.16. Reinforcement learning problem. Markov decision making process.

2.17. Reinforcement learning and the Q-learning algorithm.

2.18. Gradients by strategies and the REINFORCE algorithm.