AIE 1

Machine Learning Applications

The course focuses on basics of machine learning methods and algorithms. Finding patterns in data becomes an important part of data processing tasks essential for improvements of industrial and business activities. In the course, principles of data processing processes are presented and variety of supervised and unsupervised learning algorithms are introduced. It is shown how to evaluate and compare performance of developed models, and how diverse algorithms could lead to differences in the obtained results. Additionally, the course addresses an issue of utilization of open source machine learning tools and deployment of developed models.

In particular, the covered topics include: linear and logistic regression; support vector machines; decision tress; neural networks; and different clustering techniques.

Linear Regression I
- Cost Function
- Gradient Descent
- Performance Indexes, Testing/Training
- Library: scikit-learn

Linear Regression II
- Performance Evaluation
- Multi-feature Regression
- Normalization and Standardization
- Polynomial Regression

Logistic Regression
- Classification
- Performance Matrix
- Binary and Multi-class Classification
- Regularization

Support Vector Machine
- Maximizing the Margin
- Linear Separable using Hard and Soft Margins
- Non-linear Separable Data and use of Kernels
- SVM for regression

Decision Trees
- Introduction to Decision Tree
- Metrics used to Create Decision Tree (Gini index and Entropy)
- Decision Tree for Classification and Regression
- Ensembles and Different Techniques using Ensembles (Random Forest, Bagging and Boosting)

Neural Networks
- Introduction
- Simple Logic with Neural Networks and Hand-picked Parameters
- Simple form of Neural Networks (multilayer perceptron network)
- Mathematical Representation of Neural Network for Learning Parameters
- Linear softmax Classifier with Simple Neural Network

Unsupervised Learning I
- Clustering: k-means algorithm, Elbow Method
- Hierarchical Clustering, Dendrogram

Unsupervised Learning II
- Other Clustering Methods: Implementation and Comparison
- Dimensionality Reduction (Principal Component Analysis)
- Anomaly Detection

Deployment (AutoML tools)
- Open-Source Solutions

AIE 2

Applications with Deep and Graphical Networks

The course provides basic knowledge in the area of deep learning which is a part of machine learning that centers around recently developed architectures of neural networks.

It starts with basic concepts of feed-forward neural network, introduces Convolutional Neural Networks (CNNs) as main building blocks of modern image processing neural networks and explains some of them. Further, it moves on to Autoencoders and Generative Adversarial Networks (GANs) showing their variations and applications. The course also covers Recurrent Neural Network (RNN) and its newer models for variety of language processing applications.

Consequently, the course tries to show how image and natural language processing algorithms can be combined.

In particular, the course includes: modern CNN architectures, such as VGG, AlexNet, ResNet; Autoencoders and their applications; GANs together with DCGANs, Pix2Pix; transfer learning; data augmentation; Recurrent Neural Networks (RNNs) and their modifications such as LSTM and GRU; as well as Sequence-to-Sequence models.

Deep Learning Introduction

- Simple feed forward Neural Network (NN)
- Forward and Backpropagation
- End-to-End networks, non-linearities, and initialization
- Deep NN
- Regularization, bias and variance in NNs
- Optimization methods, gradient checking and batch normalization

Image Processing I

- Convolution analysis and CNNs
- Deep CNNs and different CNN architectures
- Image classification task
- Transfer learning and data augmentation
- Object detection and segmentation

Image Processing II

- Adversarial examples and attacks
- One-shot learning
- Generative networks
- Hands-on application, such as satellite image analysis

Natural Language Processing

- Word vectors, such as Word2Vec and GloVe
- Language modelling methods
- RNN and its updated versions
- Sentiment Analysis and text generation
- Machine translation and Seq2Seq models
- Attention and its applications

Mixed Applications

- Transfer learning
- Mixed applications of NLP and vision, such as caption generation for images
- Tabular data analysis
- Real-time data series
- Real-time object detection in videos

Practical Tips and Deployment

- Tips and Tricks
- TensorBoard, ML pipeline, and deployment

AIE 3

Reinforcement Learning Applications

In this course, basic elements of reinforcement learning are explained. Reinforcement learning is associated with learning rules or strategies, based on the past experience, which allow agents to achieve their goals or maximize their rewards. Applications of reinforcement learning can be found in control related problems, such as, robotics, gaming, for example AlphaGo, OpenAI Five, natural language processing, for example automatic dialogue systems, and others.

The course covers basics, theoretical principles, and common approaches to reinforcement learning while focusing on practical applications.

In particular, the topics are: Markov Decision Process; dynamic programming; Monte-Carlo techniques; learning algorithms with on- and off-policy as well as Temporal Difference control; function approximation; deep Q-learning; and policy gradient methods.

Introduction to Reinforcement Learning

- The background and basic elements in Reinforcement Learning
- A toolkit for reinforcement learning algorithms: OpenAI Gym

Markov Decision Processes

- The agent-environment interface
- Markov Process
- Markov Reward Process
- Bellman equation
- Value Functions, Action-Value Functions, and Policy Functions

Dynamic Programming

- Basic idea of dynamic programming
- Iterative Policy Evaluation and its proof
- Policy Improvement
- Value iteration

Model-Free prediction and control with Monte Carlo (MC)

- Difference between prediction and control
- Monte Carlo methods for prediction
- On-policy first-visit Monte Carlo control algorithm
- Off-policy Monte Carlo control algorithm
- Weighted Importance Sampling
- Benefits of Monte Carlo algorithms

Model-Free prediction and control with Temporal Difference (TD)

- Temporal Difference for prediction
- SARSA for on-policy control
- Comparison between TD algorithms and MC
- off-policy Monte Carlo control algorithm
- Unify MC with TD
- backward and forward view of TD-Lambda

Function Approximation

- Motivation of Function Approximation
- Usage of Function Approximation
- Convergence properties of function approximators
- Batch training using experience replay

Policy Gradient

- Finite Difference Policy Gradient
- Monte-Carlo Policy Gradient
- Actor-Critic Policy Gradient

Integrating Learning and Planning

- Model-based RL
- Integrated Architectures
- Simulation-based Search