

שם הקורס:

Advanced Data Science
מדעי הנתונים למתקדמים

מספר הקורס: 236605

סמסטר: חורף תשע"ז

מרצה:	קירה רדינסקי
שעות הרצאה:	09:30-11:30 יום א'
שעת תרגול:	11:30-12:30 יום א'
דרישות קדם:	<ul style="list-style-type: none">הסתברות מ 094412אלגוריתמים 234247מבוא לתכנות מערכות 234112
אתר הקורס:	

תאור הקורס

Data science waves have brought change in multiple disciplines from science and technology through business and society, hitting with full force in the last couple of years with the re-introduction of Deep Learning paradigm. With the increasing amount of data being created in the world through the development of the web, social networks and internet commerce - the need for automated methods for data analysis is growing rapidly.

In this course we will deep dive to the cutting-edge research in NLP and related disciplines and explore the state of the art methods in those fields. We will explore practical algorithms that have been used to solve key problems in data mining and can be applied successfully to even the largest datasets.

During the course we will cover the theory and practice of decision trees, SVM, graphical models, neural network models (including: word vector representations, window-based neural networks, recurrent neural networks, long-short-term-memory models, recursive neural networks, and convolutional neural networks) and explore other data mining applications such as Clustering, Massive Graphs Mining, and Recommender Systems.

The goal of the course to make you the best data-scientist and practitioner wherever you go next.

The course will be based on a final project done on a real problem from the industry in a competition mode.

The course is aimed at graduate students with some basic knowledge of machine learning (and possibly but not necessarily of deep learning) and wish to learn more about this rapidly growing field of research.

Notice: The course will accept students based on manual registration.

Intro to ML + Decision Trees	Machine learning examples Well defined machine learning problem Decision tree learning The big picture Overfitting Random variables, probabilities
Probability and Estimation Naive Bayes Gaussian Naive Bayes Logistic Regression Linear Regression Practical Issues: Feature selection Overfitting Bias-Variance tradeoff	Bayes rule MLE MAP Conditional independence Multinomial Naive Bayes Gaussian Bayes classifiers Document classification Brain image classification Form of decision surfaces Naive Bayes - the big picture Logistic Regression: Maximizing conditional likelihood Gradient ascent as a general learning/optimization method Generative/Discriminative models minimizing squared error and maximizing data likelihood bias-variance decomposition regularization
Graphical models	Bayes nets representing joint distributions with conditional independence assumptions D-separation and Conditional Independence Inference Learning from fully observed data Learning from partially observed data EM Mixture of Gaussians clustering Learning Bayes Net structure - Chow Liu Regression: Primal and Dual forms Kernels and Kernel Regression SVMs Maximizing the margin Noise and soft margin SVM's PAC learning and SVM's Hinge loss, log loss, 0-1 loss
Perceptron + SVM	
Advanced Machine Learning Topics: Active Learning Reinforcement Learning Co-Training Boosting Natural Language Processing	
Deep Learning 1	Word Vector representations: word2vec, GloVe, language models, softmax, single layer networks Neural Networks and backpropagation -- for named entity recognition
Deep Learning 2	Recurrent neural networks for language modeling • GRUs and LSTMs -- for machine translation Convolutional neural networks -- for sentence classification Similar Items, Distance Measures, Locality-Sensitive Hashing, Similarity-Preserving Summaries - Text similarity measures and text processing
Frequent Itemsets Mining Finding Similar Items (Locality-Sensitive Hashing) Clustering Dimensionality Reduction Recommender Systems Multi Arm Bandits Massive Graphs Mining	
Link Analysis I: PageRank Link Analysis II: Link Spam, HITS	Massive Graph Mining Analysis - Social Networks, community detection, triangles, small world

רשימת ספרות

- Machine Learning, Tom Mitchell.
- Pattern Recognition and Machine Learning, Christopher Bishop.
- The Elements of Statistical Learning: Data Mining, Inference and Prediction, Trevor Hastie, Robert Tibshirani, Jerome Friedman.