MATH-825 Mathematics of Data Science

Credit Hours: 3-0

Prerequisite: None

Objectives and Goals: This course explores the mathematical foundations of modern data science and machine learning. It covers core models (regression, SVMs, neural networks), geometric concepts (subspaces, SVD), high-dimensional phenomena (concentration of measure, curse of dimensionality), and algorithmic strategies (gradient descent, kernel methods). Emphasis is placed on mathematical reasoning, theoretical derivations, and practical implications of algorithms.

Detailed Course Contents: What is Data (Science)?, Affine Linear, Polynomial and Logistic Regression, k-Nearest Neighbors, Clustering, Graph Clustering, Best-Fit Subspaces, Singular Value Decomposition, Curse and Blessing of High Dimensionality, Concentration of Measure, Gaussian Random Vectors in High Dimensions, Dimensionality Reduction àla Johnson-Lindenstrauss, Separation and Fitting of High-Dimensional Gaussians, Perceptron, Support Vector Machines, Kernel Method, Neural Networks, Gradient Descent for Convex Functions, Selected Results of Probability Theory.

Course Outcomes: Upon successful completion, students will be able to understand and implement mathematically grounded machine learning models; apply concepts of linear algebra, probability, and optimization to data science; analyze high-dimensional data through subspace methods, Gaussian models, and dimensionality reduction; evaluate the theoretical guarantees and geometric properties behind machine learning algorithms; and build foundational knowledge to transition into deep learning and advanced AI methods.

Text Books:

1. Wegner, S. A. (2024). *Mathematical Introduction to Data Science*. Springer.

Reference Books:

1. Hastie, Tibshirani, Friedman (2015), *The Elements of Statistical Learning,* Springer
2. Shalev-Shwartz, S., & Ben-David, S. (2014). *Understanding machine learning: From theory to algorithms*. Cambridge university press.
3. Bishop, C. M., & Nasrabadi, N. M. (2006). *Pattern recognition and machine learning* (Vol. 4, No. 4, p. 738). New York: springer.
4. Hastie, T. (2009). The elements of statistical learning: data mining, inference, and prediction.
5. Bandeira, A. S. (2015). Ten lectures and forty-two open problems in the mathematics of data science. *Lecture notes*.

Weekly Breakdown		
Week	**Chapter/Section**	**Topic**
1	Ch 1, 2.1	What is Data (Science)? Affine Linear
2	2.2-2.4	Affine Linear, Polynomial and Logistic Regression
3	Ch 3	k-Nearest Neighbours
4	Ch 4	Clustering
5	Ch 5	Graph Clustering
6	Ch 6	Best-Fit Subspaces
7	Ch 7	Singular Value Decomposition
8	Ch 8	Curse and Blessing of High Dimensionality
9	**Mid Semester Exam**	
10	Ch 9	Concentration of Measure
11	Ch 10	Gaussian Random Vectors in High Dimensions

12	Ch 11	Dimensionality Reduction à la Johnson-Lindenstrauss
13	Ch 12, Ch 13	Separation and Fitting of High-Dimensional Gaussians, Perceptron
14	Ch 14	Support Vector Machines
15	Ch 15	Kernel Method
16	Ch 16	Neural Networks
17	Ch 17	Gradient Descent for Convex Functions
18	**Mid Semester Exam**	