AMCS 215
Mathematical Foundations of Machine Learning
Course Syllabus, Fall 2023

Course Description.
The course introduces mathematical foundations underlying neural networks and modern algorithms for regression, classification, clustering, and dimension reduction, and related supervised and unsupervised learning problems in data-rich settings. These mathematical tools, needed to understand machine learning algorithms, are traditionally taught in disparate courses, making it hard for ML students to efficiently learn them. The course bridges a gap between mathematical and machine learning courses, introducing the mathematical concepts with a minimum of prerequisites and in the context of machine learning and data science applications. The goal is to build intuition into these mathematical concepts and practical experience with applying them. Numerical computations and applications with real data will accompany the theory.

Logistics.
Class meets on Sun and Wed, 10:00–11:30am in Room 9-2322 (Lecture Hall 1). We will also have occasional sessions on Tuesdays for practical labs and to make up some travel-related absences. The class is taught by Prof. George Turkiyyah, office 1-3148, phone 808-0414, email `george.turkiyyah@kaust.edu.sa`. OH: Sun and Wed 1-2pm, or by appointment. Course TAs will be announced on Blackboard shortly.

Learning Outcomes.
The goals of the course are to introduce students to the mathematical tools used in supervised and unsupervised learning and bridge the gap between mathematical algorithms and data science and machine learning applications. At the end of the course, students are expected to develop:

- a working knowledge of the mathematics used in Data Science and Machine Learning;
- the skills needed be able to cast machine learning problems in a mathematical framework; and
- a sufficient understanding of the methods, applications, and limitations to be able to read the literature.

Tentative Schedule.

- Introduction. Machine learning problems. Collecting and cleaning data sets (1 week)
- Foundational ideas in linear algebra, optimization, and basic statistics (5 weeks)
- Regression. Logistic regression. Least squares. Ridge and Lasso regression. (1 week)
- Neural Networks. Multilayer perceptrons. Activation functions. Deep networks. Backpropagation. Training. Convolutional networks. Recurrent networks. Transformers. (5 weeks)
- Dimensionality reduction. Rank-k approximation. Eckart-Young bounds. Randomized algorithms for SVD. PCA. Matrix completion problems. Randomized projection methods. Geometry of high dimensional spaces. JL Lemma. (1 week)
- Clustering. K-means. Clustering with Mixture models. Voronoi diagrams. kD trees. Nearest neighbor classification. (1 week)
- Graph representations. Graph Laplacian. Fiedler vectors. Spectral clustering. Clustering with multiple eigenvectors. (1 week)

Grading.
The course grade will be based on:

- Homework (50%). Homework will be assigned weekly—every Sunday due the following Sunday. Keeping up with the homework is the best way to make progress on the material.
- Two exams: a midterm (25%) around week 8 and a final (25%).

Resources.
Resources for the course include:

- Probabilistic Machine Learning, Kevin Murphy, MIT Press, 2022
- Mathematics for Machine Learning, Deisenroth, Faisal, Ong, Cambridge Press, 2020
- Linear Algebra and Optimization for Machine Learning, Charu Aggrawal, Springer, 2020
- Optimization for Data Analysis, S. Wright, B. Recht, Cambridge University Press, 2022