MAE298 Machine Learning

Instructor:	Prof. Iman Soltani – isoltani@ucdavis.edu
Prereqs:	Probability and random variables, linear algebra (understanding of fundamentals) Signals and systems (Desired), Python (Ability to program is python is a must, all examples and assignments are in Python)
Textbook:	No requirement, Class notes only
Refs:	• Trevor Hastie , Jerome Friedman, Robert Tibshirani, The Elements of Statistical Learning Data Mining, Inference, and Prediction, Springer Series in Statistics (SSS), 2001. • Ian Goodfellow, Yoshua Bengio, Aaron Courville, Deep Learning, MIT Press, 2016 • Richard S. Sutton, Andrew G. Barto, Reinforcement Learning: An Introduction, MIT Press, 2018 • Kevin Patrick Murphy, Probabilistic Machine Learning: An Introduction, MIT Press, 2022 • Kevin Patrick Murphy, Probabilistic Machine Learning: Advanced Topics, MIT Press, 2023
Objectives:	The objective is to provide students with the fundamental concepts of ML and familiarize them with the latest tools and techniques so they can identify opportunities in this domain and use machine learning and deep learning in their research and future career.
Topics:	1. Fundamentals of machine learning (2-3 Weeks) • Overview of supervised, unsupervised, semi-supervised, self-supervised and reinforcement learning. • Deeper discussions on supervised learning (regression, classification, generalized linear models, support vector machines, decision tree/random forest). • Semi-supervised learning (Transductive SVM) • Unsupervised learning (EM, k-means, hierarchical, DBSCAN, dimensionality reduction) • Model training and evaluation (Train/test splits, cross-validation, bias-variance tradeoff, precision, recall, ROC) • Backpropagation, stochastic gradient descent 2. Neural networks and deep learning (2 Weeks) • Multilayer perceptron • Convolutional networks • Why deep? • Recurrent neural networks • Exploding and vanishing gradients, LSTMs • Regularization techniques (L1, L2, dropout, batch normalization, early stopping, augmentation, label smoothing) • Transfer learning • Deeper dive into self-supervised learning 3. Reinforcement learning (2 weeks) • Overview of fundamentals (MDPs, dynamic programming, Monte Carlo methods) • Temporal difference, Policy iteration and value iteration, Q-learning

	• Deep RL (Deep Q-network, experience replay) • Policy gradient, actor-critic 4. Brief intro to advanced topics (2 week) • Generative networks (Variational autoencoders, generative adversarial networks) • Attention mechanisms • Transformers and Language models 5. Research examples (1 Week) • Autonomous driving (Example: Meta learning for naturalistic driving) • Industrial diagnosis (Example: Automated health monitoring of vehicle transmission) • Healthcare (Example: Automated diagnosis of ophthalmic pathologies) • Robotics and causality (Example: Bimanual dexterity) 6. Special topics for Mechanical and Aerospace Engineers (0.5 week) • Physics informed deep learning