

CS722/822 Machine Learning

Fall 2025

Instructor

Dr. Mahmoud Nazzal, Assistant Professor, CS Department
Office: Room 3210, ENGR & COMP SCI BLDG, Norfolk, VA 23529, USA
Email: mahmoud [dot] nazzal [at] odu [dot] edu
Personal site: fs.wp.odu.edu/mnazzal
Research profile: mahmoudkanazzal.github.io
Phone:
Class times: 3 credit hours. T: 06:00 pm-08:40 pm
Office Hours: T: 05:00 pm-06:00 pm or by appointment
Class location:

Teaching Assistant

Prerequisites

Linear algebra, basic probability theory, basic data structures and algorithms, computer programming (Python)

Textbook and References

- Primary Course Material: **Lecture slides and instructor-prepared notes**.
- Textbook (assumed):
Aurélien Géron, Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow (3rd Edition), O'Reilly Media, 2022. ISBN: 9781098122478

Optional References (for depth and additional reading)

- *Christopher M. Bishop, Pattern Recognition and Machine Learning*, Springer, 2006. ISBN: 9780387310732 (*For mathematical foundations and probabilistic modeling.*)
- *Kevin P. Murphy, Machine Learning: A Probabilistic Perspective*, MIT Press, 2012. ISBN: 9780262018029
(*Advanced theoretical treatment with a focus on probabilistic approaches.*)

Brief Course Description

This graduate course is a rigorous, hands-on introduction to **classical machine learning**. We cover supervised learning (linear & polynomial regression; logistic/softmax classification; k-NN; LDA/QDA; **SVMs**; **decision trees**; **ensemble methods**) with an emphasis on evaluation, cross-validation, calibration, and error analysis. Unsupervised learning includes **clustering** (k-means, hierarchical, DBSCAN, Gaussian mixtures/anomaly detection) and **dimensionality reduction** (PCA, randomized/incremental PCA, LLE, random projections). **Neural networks are introduced only as perceptron/MLP classifiers** to place them within the ML toolbox. A late-semester **survey** highlights modern directions (deep learning, self-/semi-supervised learning, reinforcement learning, federated/privacy-preserving learning) and points students to dedicated courses. The course includes one curated paper-review session (after instructor approval), a midterm, a team project with a conference-style presentation, and a comprehensive final exam.

Late Submission Policy

Homework assignments submitted after the deadline will incur a 10% grade deduction per late day. Submissions more than 5 days late will receive a grade of zero, unless prior arrangements have been approved by the instructor.

Course Learning Objectives

Upon successful completion of this course, students will be able to:

1. **Explain** fundamental ML concepts, problem types, common pitfalls (e.g., leakage), and the end-to-end pipeline from data prep to evaluation and deployment.
2. **Implement and evaluate** classical supervised models (linear/polynomial regression; logistic/softmax; k-NN; LDA/QDA; SVMs; decision trees; ensembles) using sound metrics and scikit-learn pipelines.
3. **Perform model selection** with cross-validation, handle class imbalance, conduct error analysis, and apply basic probability calibration.
4. **Apply unsupervised methods**, including clustering (k-means, hierarchical, DBSCAN, Gaussian mixtures) and dimensionality reduction (PCA, randomized/incremental PCA, LLE, random projections), and **assess** results appropriately.
5. **Explain ANN fundamentals** and **implement a simple MLP classifier**, selecting appropriate activations/losses and basic regularization/early stopping.
6. **Design reproducible ML experiments**, including stratified splits, hyperparameter search, ablations, and documentation/governance practices.
7. **Read and critique** a contemporary ML paper and **present** findings succinctly.
8. **Plan and execute** a team ML project: define the problem, build baselines and improved models, evaluate rigorously, and communicate results.
9. **Describe at a high level** modern directions beyond classical ML (deep learning, self-/semi-supervised learning, reinforcement learning, federated/privacy-preserving learning) and identify when specialized courses/tools are appropriate.

Grading Policy

Item	Grade Allocated
Midterm	20
Final Exam	30
Homework Assignments	20
Term Project*	30

*The project grade entails the grade of paper and proposal presentation, project presentation, and project report.

Weekly Topics

Week	Topic	Materials
1 Aug 26, 2025	Introduction: What is ML?, The Big Picture, Why use ML?, Why We Still Learn Classical ML in 2025, Applications, Types of ML, Main Challenges to ML, Testing and Validating, The end-to-end ML pipeline: (Data Mining, Data Visualization, Data Preparation, Model Selection, Fine-tuning, Launching, and Maintenance),	Géron 1, 2
2 Sep 2, 2025	Preliminaries: Review of linear algebra, probability theory, and optimization theory.	
3 Sep 9, 2025	Regression: Linear Regression, Gradient Descent, Polynomial Regression, Learning Curves, Overfitting & Generalization, Regularization Techniques (Ridge, Lasso, Elastic Net, Early Stopping)	Géron 3
4 Sep 16, 2025	Classification I: Foundational and Classical Methods MNIST dataset, <i>overview of classification tasks</i> , training a binary classifier, logistic regression (estimating probabilities, training and cost function, decision boundaries), connection to linear regression as a special case , performance measures (accuracy, precision/recall, F1, ROC/AUC), error analysis (confusion matrix, false positives/negatives), <i>evaluation pipeline (stratified train/validation/test split, cross-validation, leakage check)</i> , KNN (intuition, distance metrics, feature scaling), linear discriminant analysis (generative vs discriminative view), <i>limitations of linear models (non-linear boundaries, XOR example)</i> , multiclass classification (one-vs-rest, one-vs-one), softmax regression, multilabel and multioutput classification, logistic regression – in depth (interpretability, odds ratios)	Géron 3, 10
5 Sep 23, 2025	Classification II: Neural Network-Based Methods motivation (limits of linear models, XOR revisited), introduction to artificial neural networks (from biological to artificial neurons, logical computations with neurons), perceptron (capabilities and limits), multilayer perceptron (hidden layers, non-linear decision boundaries), universal approximation perspective (intuition) , backpropagation (high-level intuition, gradient flow), activation functions (sigmoid, tanh, ReLU, softmax), loss functions (log-loss for binary, cross-entropy for multiclass, MSE vs CE discussion), regularization in neural nets (weight decay, dropout – overview) , wrap-up (neural networks as generalization of logistic regression and foundation for deep learning)	
6 Sep 30, 2025	Unsupervised Learning I – Clustering I : (overview, hierarchical clustering, clustering evaluation, silhouette score, spectral clustering), Clustering Algorithms: k-means and DBSCAN: (k-means, k-means++ initialization, Limits of k-means, Using Clustering for Image Segmentation, Using Clustering for Semi-Supervised Learning, feature scaling/standardization, DBSCAN,	Géron 9

Week	Topic	Materials
	DBSCAN parameter selection (eps, min_samples), Other Clustering Algorithms),	
7 Oct 7, 2025	Unsupervised Learning II – Clustering II, Data reduction, and outlier detection: Gaussian Mixtures: (Using Gaussian Mixtures for Anomaly Detection, Selecting the Number of Clusters, Bayesian Gaussian Mixture Models), Anomaly and Novelty Detection, Data representation and compression (Intro)	
8 Oct 14, 2025	NO CLASS (Fall Holiday Sat–Tue)	
9 Oct 21, 2025	Project Proposal and Base Paper Presentation: student teams present a paper of their choice after instructor approval, along with project proposal.	
10 Oct 28, 2025	Midterm Exam (60 minutes) Decision Trees: Training and Visualizing a Decision Tree, Making Predictions, Estimating Class Probabilities, The CART Training Algorithm, Computational Complexity, Gini Impurity or Entropy? Regularization Hyperparameters, Regression, Sensitivity to Axis Orientation, Decision Trees Have a High Variance	Géron 6
11 Nov 4, 2025	NO CLASS (Election Day)	
12 Nov 11, 2025	Support Vector Machines (SVMs): Linear SVM Classification, soft-margin & hinge loss (C, slack variables), Under the Hood of Linear SVM Classifiers, The Dual Problem, Kernelized SVMs, Nonlinear SVM Classification, multiclass strategies (one-vs-rest, one-vs-one), feature scaling, SVM Regression	Géron 8
13 Nov 18, 2025	Ensemble Learning and Random Forests: (Voting Classifiers, Bagging and Pasting: (Bagging and Pasting in Scikit-Learn, Out-of-Bag Evaluation, Random Patches and Random Subspaces), Random Forests: (Extra-Trees, Feature Importance), Boosting: (AdaBoost, Gradient Boosting, Histogram-Based Gradient Boosting), Stacking)	Géron 7
14 Nov 25, 2025	Dimensionality Reduction: (The Curse of Dimensionality, Main Approaches for Dimensionality Reduction: (Projection, Manifold Learning), PCA: (Preserving the Variance, Principal Components, Projecting Down to d Dimensions, Using Scikit-Learn, Explained Variance Ratio, Choosing the Right Number of Dimensions, PCA for Compression, Randomized PCA, Incremental PCA), Random Projection, Kernel PCA / LLE / t-SNE, Other Dimensionality Reduction Techniques) Dictionary Learning, Autoencoders Beyond ML's Supervised and Unsupervised Learning: (Deep Learning: (DL Concepts and Architectures, The Vanishing/Exploding Gradients Problems, Faster Optimizers, Learning Rate Scheduling, Avoiding Overfitting Through Regularization, Reusing Pretrained Layers), Self-Supervised Learning, Semi-Supervised Learning, Reinforcement Learning, Other ML Learning Paradigms (Incorporating domain knowledge,	

Week	Topic	Materials
	adaptability, structure, and privacy): (Physics-Informed Machine Learning (PIML), Logic- and Constraint-Guided Learning, Meta-Learning (Learning to Learn), Multi-Modal / Cross-Modal Learning, Process-Centric Learning Paradigms, Federated and Privacy-Preserving Learning))	
15 Dec 2, 2025	Project presentations: Student teams present their final term projects	
Final Week	A final comprehensive exam	

Student Mental Health and Wellbeing

ODU is committed to supporting your mental health. The Office of Counseling Services offers free, confidential support including virtual and in-person counseling, group sessions, and crisis services. Same-day or next-day appointments can be scheduled at odu.edu/counseling services. In case of a mental health emergency—such as thoughts of self-harm, harming others, or recent assault—call 911, or contact a crisis counsellor 24/7 at 757-683-4401 (press Option 2). You may also call the Suicide & Crisis Lifeline at 988 or text "HOME" to 741741.

Academic Integrity

Old Dominion University expects honesty in all academic work. In this course, your work and conduct must comply with the **Code of Student Conduct** (odu.edu/oscai). Academic dishonesty includes:

- **Cheating:** using unauthorized assistance, materials, study aids, or information.
- **Plagiarism:** using others' language or ideas without proper acknowledgment.
- **Fabrication:** inventing, altering, or falsifying data, citations, or information.
- **Facilitation:** helping another student commit a violation or failing to report suspected violations.

Classroom disruptions that undermine instruction are also prohibited. **Suspected violations will be reported to the Office of Student Conduct & Academic Integrity and may result in sanctions up to and including expulsion.**

Honor Code

Students must follow the ODU Honor Code for all assignments and exams. Collaboration on ideas is encouraged, but **all submitted work (code, text, analysis) must be your own**. Limited use of tools (e.g., ChatGPT) for brainstorming is acceptable; **do not submit AI-generated text/code as your own**. Suspected violations will be handled under University policy.

Drop Policy

As per University guidelines. See the University Calendar for drop dates.

Accessibility & Accommodations

ODU provides reasonable accommodations under the ADA. If you need accommodations, **obtain an accommodation letter from the Office of Educational Accessibility (OEA)** and share it with me early so we can implement the approved arrangements. If you anticipate barriers but don't yet have a letter, **contact OEA** to discuss eligibility.

- **OEA:** 1021 Student Success Center, (757) 683-4655, odu.edu/educationalaccessibility
- Accommodations begin **once** I receive your official OEA letter.

Attendance Policy

Students are expected to attend classes regularly.