

2026 Spring: CS4375-02 Introduction to Machine Learning

Instructor	Feng Chen
Office	ECSS 3.901
Number	(518) 442-4270
Email	feng.chen@utdallas.edu
Office Hour	MoWe 11:30am-12:30am

TA	(to be announced)
Office	
Email	
Office Hour	

Class Time	MoWe 8:30AM - 9:45AM
Location	ECSS 2.305

Course Description:

A course in machine learning (ML) introducing basic concepts and techniques. Topics include supervised learning (linear regression, logistic regression, artificial neural networks, support vector machines), unsupervised learning (K-Means, hierarchical clustering), dimension reduction (principal component analysis, linear discriminative analysis, t-stochastic neighbor embedding), reinforcement learning, and optimization.

Textbooks & References

There is no required textbook, but the following books may serve as useful references for different parts of the course.

- *Pattern Recognition and Machine Learning (PRML)* by Christopher M. Bishop
- *Machine Learning (ML)* by Tom Mitchell
- *Deep Learning (DL)* by Ian Goodfellow and Yoshua Bengio and Aaron Courville

Course Description:

Part	Lecture	Lecture Topics	Reading Materials
Introduction	1	Syllabus; Introduction	PRML (Chapter 1); DL (Chapter 1)
Supervised learning - 1	2	Linear Regression, Logistic Regression, Gradient Descent	Animation: Gradients and Partial Derivatives Visualization of gradient descent on the contour plot of the cost function Animation: Gradient and graphs Par I and Part II in Andrew Ng's Lecture Notes First 15 pages of Chapter 3 in PRML (Sections 3.1 and 3.2) Machine Learning Part 1 Andrew NG's short videos Section 3 Logistic Regression lecture by Tom Mitchell Parts 3.1.1 to 3.4.4 in Andrew NG's short videos
Artificial Neural Networks	3	Neurons, activation function, loss function, back propagation, fully connected neural networks, and convolutional neural networks	ML (Chapter 4) , DL (Chapter 6, 9) Parts 4.1.1 to 4.3.3 in Andrew NG's short videos Introduction to Deep Learning Part 1 Introduction to Deep Learning Part 2 Activation Functions Why Non-linear Activation Functions Train/Dev/Test Sets Parameters vs Hyperparameters Hyperparameter Tuning in Practice Nuts and Bolts of Applying Deep Learning Improving deep neural networks: hyperparameter tuning, regularization and Online demo of NN
Unsupervised Learning	4	K-Means, Gaussian Mixture Model	Tutorial on K-means and hierarchical clustering
Supervised Learning - 2	5	Support Vector Machines, Gradient Boosting	Learning to implement SVM in Python An introduction for SVM
Reinforcement Learning	6	Markov random fields, value iteration algorithm, Q-learning	Lecture by Serena Yeung Applications of Deep Reinforcement Learning Demo of Deep Reinforcement Learning
Dimension Reduction	7	Principal Component Analysis, Linear Discriminative Analysis, t-stochastic neighbor embedding	Dimensionality Reduction (PCA and LDA) with Practical Implementation Visualizing Data using t-SNE
Optimization Techniques for deep learning models	8	Mini-batch Gradient Descent, Momentum, Nesterov Momentum, AdaGrad, RMSProp, Adam	DL (Chapter 8) The Evolution of Gradient Descent How to implement linear regression using gradient descent Interpretation of bias correction in the Adam algorithm

Examinations and Grading:

- Grading Distribution:
 - Homework Assignment: 30%
 - In-Class Exercises: 10%
 - Midterm Exam: 30%
 - Final Project (at most two-member team): 30%
- Homework Requirement:
 - All assignments must have your name, student NetID, course name/number, and section number written in the first page of your report.
 - For each assignment, if you answered at least 50% of questions correctly, you are allowed to submit your corrections within the one week after the grading result is announced.
- In-Class Quizzes:
 - There will be several in-class quizzes. Each quiz has several questions. You will be provided printed exam-sheet for each quiz in class.
 - If you can answer 50% of the questions accumulated from all the quizzes, you will get the full marks for this portion.
 - For example, if we have 6 quizzes and each has 2 questions. There will be 12 questions in total. If you can answer at least 6 questions correctly in total, you will get full marks for this portion.
- Class Attendance:
 - Class attendance in person is mandatory.
 - You can miss up to 2 class sessions without penalty in your final numerical grade. However, if you miss more than 2 class sessions without a proper explanation, this may cause 20%-40% less from your final numerical grade.

Homework

- Homework Assignments
 - Assignments will be due at the beginning of the class on the due date
 - All assignments must have your name, student ID, course name/number, and section number
 - File Name:
 - If a report is required, the report name should have the format: 'HW_#_FirstName_LastName.pdf' (e.g., 'HW_1_Feng_Chen.pdf')
 - If the code files are required, store all the code files in one ZIP file and use the name format `HW_#_code_FirstName_LastName.zip` (e.g., `HW_1_code_Feng_Chen.pdf`)
 - Submit HW on the eLearning website.

- Late Homework
 - 30% penalty for the first 24-hour
 - 70% penalty for \geq 24 hour
 - Zero score for \geq 3 days

Course Project Requirement

- Timeline
 - Project presentation: The exam time during the exam week (May-12 to May-16)
 - Report submission deadline: May. 14 (Wed.)
- Project Report
 - [IEEE two-column style format](#), Times New Roman, Font 10.
 - Each team should submit a single report that should be 8 pages or more (excluding references). In each section/sub-section in the report, the contributions of individual team members should be made clear. For example, if students A and B contributed 80% and 20% in the writing of the section II "Related Work", respectively, then the following info should be provided in the section title: "Related Work (A: 80%, B: 20%)".
 - Useful links: [Writing Technical Articles](#), [Writing a Technical Report](#), [Paper Writing and Presentations](#)
 - The project report may contain the following sections:
 - Author (name, student ID, E-mail, and department)
 - Abstract
 - Introduction (including motivation, problem statement, significance of the problem)
 - Related work
 - Proposed approaches
 - System design and implementation (including architecture, datasets, major components, graphic user interface)
 - References
 - Example project reports
 - [Example project report 1](#)
 - [Example project report 2](#)
 - [Example project report 3](#)
 - [Example project report 4](#)
 - [Example project report 5](#)
- **Policy on Cheating:**

Cheating in an exam will result in an E grade for the course. Further, the students involved will be referred to the Dean's office for disciplinary action. Homework problems are meant to be individual exercises; you must do these by yourself. Any of the following actions will be considered as cheating.

 1. A solution which is identical to or nearly identical to the solution submitted by another student in the class.
 2. A solution which is identical to or nearly identical to the solution provided by the instructor in a previous offering of the course.
 3. A solution which is identical to or nearly identical to a solution available on the Internet.
 - Cheating in a homework exercise will result in the following penalty for all the students involved.
 1. The homework in which cheating occurred will be assigned a grade of ZERO.
 2. The homework in which cheating occurred will be assigned a grade of ZERO.
 - Students who cheat in two or more homework assignments will receive an E grade for the course. The names of such students will also be forwarded to the Dean's office for disciplinary action.

For a complete list of UTD policies and procedures, see [here](#).