Machine Learning 2026: Course Handout

Instructor: Prof. Udit Bhatia, Civil Engineering; Computer Science and Engineering
Tutor: Prof. Krishna Prasad Miyapuram, Cognitive and Brain Sciences; Artificial Intelligence
Semester I, 2026
ES 335 Machine Learning, UG track | ES 678 Machine Learning, PG track
Credits: 3 | 1.5 | 0 | 4

Course Description

Machine learning in 2026 is no longer only about fitting a model in a notebook. It is the discipline of building data-driven systems whose predictions, explanations, and decisions remain reliable outside the training environment. This course teaches the foundations behind such systems: data, loss, risk, optimization, generalization, evaluation, and model behavior under real-world constraints.

The course will use modern tools, including AI coding assistants, but it will not treat them as substitutes for understanding. Technologies will change quickly; the mathematical and statistical ideas used to reason about learning will remain portable across domains.

Course Objectives

By the end of the course, students should be able to:

1. Formulate a supervised or unsupervised learning problem with clear notation for data, targets, models, loss, and evaluation.
2. Derive and interpret core ML objectives, including empirical risk minimization and common regression and classification losses.
3. Diagnose overfitting, leakage, distribution shift, metric mismatch, shortcut learning, and deployment risk.
4. Implement and evaluate classical ML models, including regression, logistic regression, trees, ensembles, KNN, SVMs, PCA, and clustering.
5. Explain the core ideas behind neural networks, transformers, foundation models, agents, and modern ML systems.
6. Use AI/GenAI tools responsibly while remaining accountable for equations, code, plots, assumptions, and conclusions.

Prerequisites

Expected background:

- Basic linear algebra: vectors, matrices, projections, eigenvalues at an introductory level.
- Basic probability and statistics: random variables, expectation, variance, conditional probability, sampling.
- Basic programming: Python, notebooks, arrays, plotting, and reading documentation.
- Willingness to do derivations by hand.

Students who are rusty in any of these areas should revise them during the first two weeks.

Syllabus Outline

1. ML in 2026: past, present, systems, and failure modes.
2. What it means to learn from data: population, sample, features, labels, loss, risk.
3. Mathematical foundations: linear algebra, probability, statistics, optimization.
4. Generalization and evaluation: train/validation/test, leakage, metric choice, thresholds, costs.
5. Regression: least squares, gradient descent, regularization, diagnostics.

6. Classification: logistic regression, likelihood, cross-entropy, calibration.
7. Classical models: decision trees, ensembles, KNN, SVMs, kernels.
8. Representation and unsupervised learning: PCA, clustering, embeddings.
9. Neural networks: forward pass, backpropagation, optimization, regularization.
10. Modern ML: transformers, foundation models, agents, evaluation, deployment, accountability.

Timetable and Teaching Calendar

Weekly timetable slots:

- Lectures: F1 on Monday 11:30–12:50 and F2 on Thursday 11:30–12:50, room 10/103.
- Tutorial: B1 on Tuesday 08:30–09:50, Surendra C Sheth Learning Theatre / Jayshree S Sheth Learning Theatre.
- There are no regular ML meetings during Examination I or the mid-semester recess.

Lecture Calendar

Lec.	**Slot/date**	**Topic**	**Main learning objectives**	**Assessment connection**
1	Thu Aug 06	Machine Learning: introduction, past and present	Understand ML in 2026; distinguish tools from fundamentals; see paradoxes and course expectations.	A0 active
2	Mon Aug 10	What does it mean to learn from data?	Define population, sample, features, target, model, loss, risk, and empirical risk.	A0 preparation
3	Thu Aug 13	Linear algebra I: data as vectors and matrices	Interpret samples, features, labels, parameters, norms, dot products, and matrix notation.	Q1 window opens
4	Mon Aug 17	Linear algebra II: geometry and projection	Use projections, orthogonality, rank, and least-squares geometry to reason about models.	A1 preparation
5	Thu Aug 20	Probability and statistics I	Work with random variables, expectation, variance, conditioning, and Bayes rule.	A1 active
6	Mon Aug 24	Probability and statistics II	Connect sampling, estimation, uncertainty, likelihood, and confidence to ML claims.	Q1 possible
7	Thu Aug 27	Generalization, evaluation, and leakage I	Explain train/validation/test splits, leakage, distribution shift, and the role of held-out data.	A1 support
8	Mon Aug 31	Generalization, evaluation, and leakage II	Diagnose misleading scores, shortcut signals, metric mismatch, and real-world evaluation risk.	A1 due in tutorial
9	Thu Sep 03	Optimization I	Connect objectives, gradients, convexity, learning rates, and descent directions.	A2 released
10	Mon Sep 07	Optimization II	Study gradient descent behavior, numerical stability, scaling, and convergence diagnostics.	Q2 possible
11	Thu Sep 10	Linear regression I	Derive least squares; interpret coefficients, residuals, assumptions, and normal equations.	A2 support
12	Mon Sep 14	Linear regression II	Study diagnostics, nonlinear features, variance, error analysis, and model comparison.	Exam I preparation
13	Thu Sep 17	Regularization and feature design	Explain ridge/lasso intuition; connect capacity, bias-variance, shrinkage, and generalization.	Exam I preparation
-	Sep 18-23	Examination I period	No regular ML meetings.	Mid-semester exam
14	Thu Sep 24	Logistic regression I	Derive logistic loss from likelihood; interpret probabilities and decision boundaries.	Q3 window opens
15	Mon Sep 28	Logistic regression II	Connect cross-entropy, calibration, thresholds, odds, and classification uncertainty.	A2 due in tutorial
16	Thu Oct 01	Metrics, thresholds, and costs	Compare accuracy, precision, recall, ROC, calibration, cost-sensitive decisions, and base rates.	A3 released
17	Mon Oct 05	Validation and experimental design	Design honest experiments; avoid leakage; choose baselines, controls, and ablations.	Q3 possible

Lec.	Slot/date	Topic	Main learning objectives	Assessment connection
18	Thu Oct 08	Decision trees I	Explain splits, impurity, depth, greedy learning, and interpretability.	A3 support
19	Mon Oct 12	Decision trees II	Study pruning, instability, variance, categorical features, and missing-data issues.	A3 support
20	Thu Oct 15	Ensembles	Understand bagging, random forests, boosting, variance reduction, and robustness.	A3/A4 due; A5 released
-	Oct 17-25	Mid-semester recess	No regular ML meetings.	Dussehra falls during recess
21	Mon Oct 26	KNN and instance-based learning	Reason about neighborhoods, distance metrics, scaling, and curse of dimensionality.	Project released in tutorial
22	Thu Oct 29	SVMs and kernels	Understand margins, hinge loss, kernels, support vectors, and feature spaces.	Q4 possible
23	Mon Nov 02	PCA and representation learning	Derive PCA intuition; use eigenvectors/SVD for dimensionality reduction and denoising.	Project proposal due in tutorial
24	Thu Nov 05	Clustering and unsupervised learning	Compare k-means, hierarchical clustering, validation, and interpretation risks.	A4 support
25	Mon Nov 09	Neural networks I	Explain neurons, layers, activations, forward pass, losses, and representation learning.	A5 due in tutorial
26	Thu Nov 12	Neural networks II	Explain backpropagation, optimization pathologies, regularization, and double descent.	Q5 possible
27	Mon Nov 16	Transformers, foundation models, agents	Understand attention, pretraining, prompting, tool use, and evaluation challenges.	Project final due in tutorial
28	Thu Nov 19	ML systems, deployment, and accountability	Connect monitoring, drift, reproducibility, human oversight, and responsible deployment.	Q6 possible; final synthesis
-	Nov 20-28	Examination II period	End-semester examination window.	End-semester exam

Tutorial Calendar

Tut.	Date	Tutorial focus	Output or assessment connection
0	Tue Aug 04	Course launch; notebook protocol; A0 mathematical readiness; notation and diagnostic problems.	A0 released
1	Tue Aug 11	A0 discussion; scalar/vector/matrix notation; dimensions; projections; hand derivation practice.	A0 due
2	Tue Aug 18	Probability and statistics problem clinic; Bayes rule, expectation, variance, sampling, and units.	A1 released
3	Tue Aug 25	Data audit and leakage clinic; identifying sampling assumptions, shortcuts, and measurement issues.	A1 support
4	Tue Sep 01	Optimization clinic; gradients, step sizes, numerical checks, and loss-surface sketches.	A1 due
5	Tue Sep 08	Regression derivation clinic; normal equations, gradient descent, residual plots, and diagnostics.	A2 support
6	Tue Sep 15	Regularization and Exam I review; bias-variance, ridge/lasso intuition, and closed-book practice.	Exam I preparation
-	Tue Sep 22	No tutorial: Examination I period.	-
7	Tue Sep 29	Logistic regression clinic; likelihood, cross-entropy, calibration, and threshold reasoning.	A2 due; A4 released
8	Tue Oct 06	Metrics and validation clinic; confusion matrices, ROC/PR tradeoffs, baselines, and ablations.	A3 support
9	Tue Oct 13	Trees and ensembles clinic; impurity calculations, pruning, random forests, and boosting logic.	A3/A4 support
-	Tue Oct 20	No tutorial: mid-semester recess.	-
10	Tue Oct 27	KNN/SVM clinic; scaling, distance choices, margins, kernels, and support-vector reasoning.	Project released
11	Tue Nov 03	PCA, clustering, and project proposal clinic; SVD intuition, k-means failure modes, and baselines.	Project proposal due
12	Tue Nov 10	Neural-network clinic; backpropagation, parameter counting, regularization, and debugging.	A5 due
13	Tue Nov 17	Project viva preparation and final review; assumptions, baselines, failure modes, and accountability.	Final project due

Assignment Calendar

Item	Release	Due	Scope
A0	Tue Aug 04	Tue Aug 11	Graded mathematical readiness assignment: notation, dimensions, linear algebra, calculus, probability, statistics, and opening lecture vocabulary.
A1	Tue Aug 18	Tue Sep 01	Data audit, visualization, probability/statistics, leakage or paradox analysis.
A2	Thu Sep 03	Tue Sep 29	Optimization, regression, regularization, diagnostics, and interpretation.
A3	Thu Oct 01	Thu Oct 15	Classification, metrics, thresholds, validation, trees, and failure-mode audit.
A4	Tue Sep 29	Thu Oct 15	Theory, reproducibility, and failure modes.
A5	Thu Oct 15	Tue Nov 10	Classical models and representation: ensembles, KNN/SVM, PCA, and clustering.
Project	Tue Oct 27	Tue Nov 17	End-to-end ML project with proposal due Tue Nov 03. PG students include addendum.

Quizzes and Exams

- Expect quizzes at any time. Quizzes may be surprise or announced.
- No phones, laptops, tablets, smart watches, or electronic devices are allowed during quizzes unless explicitly announced.
- Quizzes may be closed book, closed notes, or open notebook if explicitly announced.
- If an open-notebook quiz is announced, only the dedicated course notebook is allowed.
- Approximately six quizzes may be conducted; if more than five are conducted, the best five may be counted at the instructor's discretion.
- Mid-semester and end-semester exams are closed book and closed notes unless explicitly stated otherwise.

Grading

ES 335: UG Track

Component	Weight
Quizzes	10%
Assignments A0-A5	30%
Mid-semester exam	20%
End-semester exam	25%
Course project	15%

ES 678: PG Track

Component	Weight
Quizzes	10%
Assignments A0-A5	30%
Mid-semester exam	20%
End-semester exam	25%
Course project with PG addendum	15%

A0 is graded and included in the assignment component. All assignment questions are common to ES 335 and ES 678 students. PG students receive additional assessed depth through PG-specific exam questions and a research/depth component in the project.

Submission Policies

1. Every assignment must be submitted as a single PDF file.
2. File name format: <roll_number>_<first_name>_<last_name>.pdf.
3. The PDF must be self-contained: answers, derivations, plots, tables, assumptions, and conclusions should be in the main report.
4. Code should be included as an annexure or submitted separately only when explicitly requested.
5. Evaluation will focus on correctness, clarity, reasoning, and learning rather than length.
6. AI/GenAI tools are permitted for assignments, but students are accountable for every final answer, equation, plot, claim, and line of code.
7. If AI tools are used materially, include a brief note describing where they were used and what was independently checked.
8. A viva on any assignment may happen at any time as deemed fit by the instructor.
9. Unless an extension is announced or approved in advance, deadlines are hard. Late submissions may be penalized or not accepted after solutions, rubrics, or viva sampling begin.

Classroom Conduct

1. No phones or laptops are allowed in class unless explicitly requested for a tutorial or coding activity.
2. Class doors close after 5 minutes. There is no attendance credit as such, but no credit or make-up will be given if a surprise quiz happens after a student is late or absent.
3. Students must maintain a separate course notebook and bring it to every class.
4. Only the dedicated course notebook will be allowed for open-notebook quizzes or exams, if any are announced.
5. The notebook should contain definitions, derivations, mistakes, questions, and reflections. It is part of how students learn the subject.

Project Policy

The project must be an end-to-end ML study, not just a model-training exercise. A good project includes:

- a clear problem statement and data description;
- a baseline that is simple and meaningful;
- a justified model and evaluation protocol;
- error analysis and at least one failure-mode discussion;
- reproducible code and a concise report;
- clear separation between what the model predicts and what decision the prediction supports.

PG students must add a research/depth component, such as a paper discussion, nontrivial baseline, derivation, ablation, calibration analysis, sensitivity analysis, or deployment-risk audit.

Academic Integrity and AI Accountability

Collaboration is encouraged at the level of discussion, debugging, and conceptual clarification. Submitted work must reflect the student's own understanding. Students may use AI systems for coding help, explanation, editing, plotting, or debugging, but they remain responsible for the final submission. During a viva, students should be able to explain any derivation, code block, plot, or modeling choice in their submission.

Central rule: tools may assist the work, but they cannot carry the accountability.

Reference Books

1. Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong, *Mathematics for Machine Learning*, Cambridge University Press.
2. Gareth James, Daniela Witten, Trevor Hastie, Robert Tibshirani, and Jonathan Taylor, *An Introduction to Statistical Learning*, Springer.
3. Christopher M. Bishop, *Pattern Recognition and Machine Learning*, Springer.
4. Trevor Hastie, Robert Tibshirani, and Jerome Friedman, *The Elements of Statistical Learning*, Springer.
5. Richard O. Duda, Peter E. Hart, and David G. Stork, *Pattern Classification*, Wiley.
6. Tom M. Mitchell, *Machine Learning*, McGraw-Hill.
7. Kevin P. Murphy, *Machine Learning: A Probabilistic Perspective*, MIT Press.
8. Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning*, MIT Press.

Acknowledgement

Some course materials are informed by Prof. Nipun Batra's Machine Learning teaching resources at IIT Gandhinagar:

- `https://github.com/nipunbatra/ml-teaching`
- `https://nipunbatra.github.io/ml-teaching/`
- `https://nipunbatra.github.io/ml2024/`