Machine Learning 2026: Assessment Structure

Planning document for the dual-listed UG/PG offering
ES 335 Machine Learning | ES 678 Machine Learning
Prof. Udit Bhatia, Civil Engineering; Computer Science and Engineering
Prof. Krishna Prasad Miyapuram, Cognitive and Brain Sciences; Artificial Intelligence
Credits: 3 | 1.5 | 0 | 4

Working Assumption

The course is dual listed with one common classroom offering.

- UG code: ES 335 Machine Learning.
- PG code: ES 678 Machine Learning.
- Shared credits: 3 | 1.5 | 0 | 4.

The course should remain a single course in lectures, tutorials, quizzes, and assignments. All assignment questions should be common to UG and PG students. PG students should have additional assessed depth through a project addendum and PG-only exam questions.

Core Principle

All students must master the same ML foundations:

- learning from data, empirical risk minimization, generalization, overfitting;
- linear algebra, probability, statistics, optimization;
- regression, classification, trees, ensembles, SVMs, PCA, clustering;
- neural networks, transformers, foundation models, evaluation, leakage, deployment, and accountability.

PG students should be assessed on deeper mathematical reasoning, paper reading, experimental design, and the ability to defend choices under real-world ambiguity.

Recommended Assessment Split

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%

This keeps the assignment stream common and comparable. The code-level distinction for PG students comes from exam questions and the project addendum, not from separate assignment questions.

Common Advanced Assignment A4

Suggested title: **Theory, Reproducibility, and Failure Modes**

Suggested timing:

- Release: Tuesday, September 29, 2026.
- Due: Thursday, October 15, 2026.
- This places it after the mid-semester exam window and before the mid-semester recess.

Recommended structure:

1. **Mathematical derivation.** Derive one ML loss from a probabilistic model. Good options include logistic loss from Bernoulli likelihood, ridge regression from MAP with Gaussian prior, or cross-entropy from maximum likelihood.
2. **Generalization or optimization argument.** Prove or carefully explain a bound, convergence claim, convexity condition, or regularization effect. The emphasis should be correctness of reasoning, not length.
3. **Reproducibility exercise.** Reproduce a small result from a paper, textbook example, or high-quality public implementation. Students must report what matched, what did not, and why.
4. **Failure-mode audit.** Analyze leakage, distribution shift, calibration, unfair subgroup performance, data quality, or metric mismatch. Students may use AI tools, but must document prompts, tool outputs used, and final human decisions.

Submission format:

- One PDF using the same roll-number naming convention as all assignments.
- Code as annexure.
- Short viva may be conducted for any submission.

No part of A4 should be marked UG-only or PG-only. If an item is mathematically deeper, it should still be assigned to everyone, with grading focused on reasoning and clearly shown work.

PG Project Addendum

The project topic can be the same as UG, but PG students must submit an additional research-style addendum.

Required PG addendum:

- 4–6 pages, included as a clearly marked section in the final project PDF.
- One paper, strong baseline, or serious technical reference must be discussed.
- Must contain at least two of the following:
  - derivation or theoretical justification;
  - ablation or sensitivity analysis;
  - uncertainty, calibration, or confidence analysis;
  - comparison against a nontrivial baseline;
  - failure-mode and data-leakage audit;
  - compute, data, and reproducibility checklist;
  - ethical or deployment-risk analysis tied to evidence.

Suggested project checkpoints:

- Project release: Tuesday, October 27, 2026.

- Proposal due: Tuesday, November 03, 2026.
- Final project due: Tuesday, November 17, 2026.

PG proposal requirement: in addition to the common proposal, PG students must identify the paper, baseline, or theoretical angle they will use for the addendum.

PG Exam Component

The mid-semester and end-semester exams should be common in structure, with PG-only questions clearly marked.

Recommended design:

- 75–80% common paper for all students.
- 20–25% PG-only analytical section for ES 678 students.
- UG students may optionally attempt PG questions for bonus or challenge credit only if desired by instructor policy.

Good PG-only exam question types:

- derive a loss, estimator, or update rule from assumptions;
- prove convexity or identify where convexity fails;
- analyze a generalization or bias-variance claim;
- compare two validation designs and identify leakage;
- reason about calibration, thresholding, or cost-sensitive decision-making;
- critique a foundation-model-assisted ML workflow.

Policy Text For Syllabus

This course is offered under separate UG and PG course codes. Lectures, tutorials, quizzes, and all assignment questions are common. Students registered under the PG code will have additional assessed depth through PG-specific questions in examinations and an additional research/depth component in the course project. These components are intended to assess mathematical maturity, research reading, reproducibility, and the ability to reason about real-world ML systems beyond implementation.