CS-440/740 | COURSE OUTLINE

# Machine Learning

| | |
|---|---|
| **Course code** | CS-440/740 (undergraduate section CS-440; graduate section CS-740) |
| **Credit hours** | 3 (2 + 1): two lecture hours and one three-hour laboratory per week |
| **Levels** | BS (final year) \| MS \| PhD — one course, three levels of assessment |
| **Prerequisites** | Multivariable calculus; linear algebra; probability and statistics; programming in Python. Artificial Intelligence (CS-430) is recommended but not required. |
| **Instructor** | Dr. Asif Mehmood, Department of Computer Science |
| **Contact** | asifnazquaidian@gmail.com |
| **Delivery** | Lectures with active learning, weekly programming laboratory, problem sets, graduate research project |
| **Textbook** | Machine Learning: Theory, algorithms and practice: from PAC learning to deep generative models, with applications to biomedical data (course textbook, October 2026), with the reference books below. |

## 1 Course description

Machine learning studies algorithms that improve with experience. This course builds the subject from its mathematical foundations to the methods that define current practice. It begins with the formal learning problem: empirical risk minimisation, PAC learning, uniform convergence, the VC dimension and the bias–complexity trade-off, together with model selection, cross-validation and regularisation. It then develops the core supervised methods with their derivations: linear and logistic regression and generalised linear models, generative classifiers, convex optimisation and stochastic gradient descent, support vector machines and kernels, and decision trees, bagging and boosting. The second half covers neural networks and backpropagation, deep learning in practice (convolutional networks, optimisation and regularisation, double descent), sequence models and transformers, unsupervised learning (clustering, mixture models and EM, PCA and ICA, embeddings), and probabilistic and generative models (Bayesian inference, Gaussian processes, variational autoencoders, diffusion models). It closes with machine learning for neural and physiological data, with emphasis on rigorous evaluation, leakage, fairness and reproducibility. The coverage and level follow Stanford CS229, MIT 6.790 (formerly 6.867) and the standard texts of Shalev-Shwartz and Ben-David, Bishop, Hastie, Tibshirani and Friedman, and Murphy; graduate students add proofs, primary literature and a research project.

## 2 Course learning outcomes (CLOs)

Outcomes are written at the cognitive levels of Bloom's revised taxonomy and mapped to the programme learning outcomes (PLOs) of the host programme (for computing degrees, the Seoul Accord graduate at-

| CLO | Students will be able to… | Bloom level | PLO |
|---|---|---|---|
| CLO1 | State the formal learning problem and apply PAC, VC and bias–complexity arguments to explain generalisation and model selection. | C3 Apply | PLO1, PLO2 |
| CLO2 | Derive, implement and tune linear, kernel, tree-based and ensemble learners with appropriate loss functions and regularisation. | C3 Apply | PLO2, PLO5 |
| CLO3 | Train and analyse neural networks, convolutional networks and transformers, including their optimisation and regularisation. | C4 Analyse | PLO2, PLO3 |
| CLO4 | Apply unsupervised, probabilistic and generative models to discover structure and quantify uncertainty in data. | C4 Analyse | PLO2, PLO3 |
| CLO5 | Evaluate machine-learning systems rigorously on biomedical data, detecting leakage, bias and distribution shift. | C5 Evaluate | PLO4, PLO6, PLO8 |
| CLO6 | Design, carry out and report an original machine-learning study with baselines, ablations and reproducible code. | C6 Create | PLO4, PLO9 |

tributes used by NCEAC).

## 3 Weekly schedule

Sixteen weeks: fourteen teaching weeks, a mid-term week and a final-examination week. Each teaching week has two lecture hours and one three-hour computational laboratory.

| Week | Topic (textbook chapter) | Laboratory | Assessment / CLO |
|---|---|---|---|
| 1 | The Learning Problem (Ch. 1) | Overfitting with polynomial regression | Diagnostic quiz |
| 2 | PAC Learning, Uniform Convergence and the VC Dimension (Ch. 2) | Generalisation gaps in practice | Problem set 1 (CLO1) |
| 3 | Bias, Variance, Model Selection and Regularisation (Ch. 3) | Cross-validation and learning curves | Programming lab 1 (CLO1, CLO2) |
| 4 | Linear Regression and Least Squares (Ch. 4) | Ridge and lasso on clinical data | Problem set 2 (CLO2) |
| 5 | Linear Classification: Perceptrons, Logistic Regression and Generative Models (Ch. 5) | Logistic regression from scratch | Programming lab 2 (CLO2) |
| 6 | Convex Optimisation and Stochastic Gradient Descent (Ch. 6) | Comparing optimisers | Problem set 3 (CLO2) |
| 7 | Support Vector Machines and Kernel Methods (Ch. 7) | Pegasos and kernel SVMs | Quiz 1 (CLO1, CLO2); project proposal (graduate) |
| 8 | **Mid-term examination (Chapters 1–7)** | – | Mid-term (CLO1, CLO2) |
| 9 | Nearest Neighbours, Decision Trees and Ensembles (Ch. 8) | Forests and boosting on clinical data | Programming lab 3 (CLO2) |

| Week | Topic (textbook chapter) | Laboratory | Assessment / CLO |
|---|---|---|---|
| 10 | Neural Networks and Backpropagation (Ch. 9) | A neural network from scratch, with a gradient check | Problem set 4 (CLO3) |
| 11 | Deep Learning in Practice: Convolutional Networks, Training and Generalisation (Ch. 10) | A convolutional network in PyTorch | Programming lab 4 (CLO3) |
| 12 | Sequence Models, Transformers and Representation Learning (Ch. 11) | An LSTM and a transformer on a memory task | Programming lab 5 (CLO3) |
| 13 | Unsupervised Learning: Clustering, Mixtures and Dimensionality Reduction (Ch. 12) | EM for a Gaussian mixture and PCA | Problem set 5 (CLO4) |
| 14 | Probabilistic and Generative Models (Ch. 13) | Gaussian-process regression and a VAE | Quiz 2 (CLO3, CLO4) |
| 15 | Machine Learning for Neural and Physiological Data (Ch. 14) | Leakage and honest evaluation | Project presentations (CLO6) |
| 16 | **Final examination (all chapters)** | – | Final (CLO1–CLO5); project report (CLO6) |

## 4 Assessment

Undergraduate (BS) and graduate (MS/PhD) students attend the same lectures but are assessed differently: graduate students complete the advanced (MS) and research (PhD) exercises and a research project.

### BS students

| | |
|---|---|
| Problem sets (5) | 15% |
| Quizzes (2) | 10% |
| Programming labs (5) | 20% |
| Mid-term examination | 20% |
| Final examination | 35% |

### MS and PhD students

| | |
|---|---|
| Problem sets with MS/PhD problems | 15% |
| Programming labs (5) | 10% |
| Critical review of two research papers | 10% |
| Mid-term examination | 15% |
| Research project (proposal, presentation, report) | 25% |
| Final examination | 25% |

### Graduate research project

Graduate students choose a project in week 7 and present it in week 15. A project asks a precise question and answers it with careful experiments (baselines, ablations, confidence intervals, held-out or external test data) or with theory. Examples: arrhythmia classification on PTB-XL with a 1-D CNN versus gradient-boosted features, with subject-wise splits and calibration analysis; decoding behaviour from Neuropixels recordings with linear, recurrent and transformer models; double descent in random-feature regression; a variational autoencoder for single-cell or spike data; detecting distribution shift between hospitals. Deliverables: a two-page proposal, a 15-minute presentation and a report in conference format (6–8 pages) with code and a data statement. PhD reports must state a hypothesis and position the result in the literature.

### Grading rubric for laboratory reports and projects

| Criterion | Excellent (4) | Satisfactory (2–3) | Needs work (0–1) |
|---|---|---|---|
| Model and method | Correct equations, algorithms and parameters, justified choices | Minor errors that do not change conclusions | Wrong model or unexplained method |
| Code and results | Runs, reproducible, figures labelled with units, baselines and error bars | Runs with small gaps | Does not run or results missing |
| Interpretation | Links results to theory, physiology or the literature | Describes results without depth | No interpretation |
| Communication | Clear structure, correct references | Understandable, some gaps | Unclear or unreferenced |

## 5 Reading list

### Course textbook

Machine Learning: Theory, algorithms and practice: from PAC learning to deep generative models, with applications to biomedical data. Dr. Asif Mehmood, October 2026.

### Reference books and papers

1. S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014 (principal text for theory).
2. C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006; K. P. Murphy, *Probabilistic Machine Learning: An Introduction*. MIT Press, 2022.
3. T. Hastie, R. Tibshirani and J. Friedman, *The Elements of Statistical Learning*, 2nd ed. Springer, 2009 (free online).
4. I. Goodfellow, Y. Bengio and A. Courville, *Deep Learning*. MIT Press, 2016; S. J. D. Prince, *Understanding Deep Learning*. MIT Press, 2023 (free online).
5. S. J. D. Prince, *Computer Vision: Models, Learning, and Inference*. Cambridge University Press, 2012; D. Kriesel, *A Brief Introduction to Neural Networks*, 2007.
6. J. Leskovec, A. Rajaraman and J. D. Ullman, *Mining of Massive Datasets*, 2nd ed. Cambridge University Press, 2014 (Chapters on clustering, dimensionality reduction and large-scale learning).
7. Primary papers named in each chapter, including Valiant (1984), Vapnik and Chervonenkis (1971), Cortes and Vapnik (1995), Breiman (2001), Rumelhart *et al.* (1986), He *et al.* (2016), Vaswani *et al.* (2017), Kingma and Welling (2014) and Ho *et al.* (2020).

## 6 Course policies

- Attendance of at least 80% of lectures and laboratories is required to sit the final examination.

- Problem sets and labs are individual work. Ideas may be discussed, but code and text must be written and understood by the student. Use of AI assistants must be acknowledged with a short statement of how they were used.
- Late submissions lose 10% per day and are not accepted after five days without a documented reason.
- Plagiarism is handled according to the institution's academic-integrity policy and the HEC plagiarism policy.
- Students with a disability or special need should contact the instructor in the first two weeks so that reasonable adjustments can be made.