

BIRLA INSTITUTE OF TECHNOLOGY & SCIENCE, PILANI Hyderabad Campus, Hyderabad Computer Sc. Department, 1st Sem 2024-2025 Course Handout (BITS F464: Machine Learning)

Date: 2nd August 2024

Course	Number
Course	Title
Instruct	or-In-Charge

: BITS F464 (MWF: 3rd hour) : Machine Learning : Chittaranjan Hota, Sr. Prof. (Email: hota[AT]hyderabad.bits-pilani.ac.in)

Scope and Objectives of the course:

This course is an undergraduate course on Machine Learning. ML is the sub-field of Artificial Intelligence. It helps engineers build automated systems that learn from experiences. It helps machines make data-driven decisions. For example, Google Maps for navigation uses the route network, real-time traffic characteristics, time of travel etc. to predict an appropriate path for you using ML algorithms. ML is a muti-disciplinary field, with roots in Computer Science, and Mathematics. ML methods are best described using linear and matrix algebra and their behaviour are best understood using the tools of probability and statistics. By integrating mathematical principles, you will learn to effectively address machine learning challenges, developing a deep understanding similar to that of professional data scientists. According to the latest estimates, 328 million terabytes of data are created daily. With this increasing amounts of data, the need for automated methods for data analysis continues to grow. The goal of this course is to develop methods that can automatically detect patterns in data, and then use the uncovered patterns to predict the future outcomes of interest. This course will cover many ML and Gen AI models and algorithms, including Linear regression, Multi-layer neural networks, Support vector machines, Bayesian networks, Gaussian mixture models, Clustering algorithms, Generative adversarial networks (GANs), RNNs, and Reinforcement learning techniques. Hands-on experience will be emphasized, allowing you to select optimal models and master the essential implementation details critical to their success. Practical sessions (Coding assignments) will involve working with real-world data, enhancing your proficiency in debugging and refining models through various ML techniques. The course objectives are the following:

- To understand various ML techniques like Model selection, Under-fitting, Over-fitting, Cross-validation, Regularization etc.
- To understand and build appropriate supervised learning algorithms for classification problems like Decision Trees, Naïve Bayes, Support vector machines (SVMs), Artificial Neural Networks etc.
- To understand and build appropriate supervised learning algorithms for regression problems like Linear regression, Polynomial regression, Ridge regression etc.
- To understand and build appropriate un-supervised learning algorithms for clustering, linear and nonlinear dimensionality reduction etc.
- To understand architectures and build Sequential and Generative AI models like RNNs, LSTMs, GANs, VAEs etc.
- To test run appropriate ML algorithms on real world and synthetic datasets and interpret their results over ML frameworks like ScikitLearn, TensorFlow, Keras, PyTorch etc.

Text Books:

- T1: Christopher Bishop: Pattern Recognition and Machine Learning, Springer-Verlag New York Inc., 2006.
- T2: Tom M. Mitchell: Machine Learning, The McGraw-Hill, Indian Edition, 2017.

Reference Books:

R1: Kevin Murphy: Machine Learning: A Probabilistic Perspective, MIT Press, 2012.

R2: Shai Shalev-Shwartz and Shai Ben-David: Understanding Machine Learning: From Theory to Algorithms, Cambridge University Press, 2014.

R3: Ethem Alpaydin: Introduction to Machine Learning, 3rd Edition, MIT Press, 2014.

R4: Aurelien Geron: Hands-on Machine Learning with Scikit-Learn, Keras, and TensorFlow, O'Reilly, 2019.

Lecture Plan:

Lect. #	Learning objectives	Topics to be covered	Chapter in the Text Book
1-2	Course Introduction	Course Administration, Motivation and ML Frameworks.	T2(1), Lecture Slides
3	Overview of ML	Supervised/Unsupervised/RL, Classification/ Regression, General Approach.	R3(1.1, 1.2)
4 - 5	Supervised Learning - I	Concept Learning: Version space and Candidate elimination algorithm.	T2(2.2, 2.5)
6 - 8	Supervised Learning - II	Decision Tree: Learning by Induction, Tree Representation, DT Learning algorithm, Quantifying uncertainty: Entropy, Information Gain, Gini Index, ID3 Vs C4.5 Vs CART, Over- fitting in DTs, Random Forest: Bagging, Boosting.	T2(3.2, 3.3, 3.4)
9 - 10	Evaluating a model	Bias, Variance, Cross-validation, Precision-Recall, ROC Curve, Out-of-Bag metric.	T1(1.3), R3(19.6, 19.7)
11 - 13	Linear Models for Regression	Linear regression, Logistic regression, Gradient Descent, GD Analysis, Stochastic Gradient Descent.	T1(3.1, 3.2), R1(8.1-3, 8.6)
14 - 15	Linear Models for Classification	Linear discriminant functions for Classification, Least squares for classification, Fischer's Linear Discriminant functions.	T1(4.1)
16 - 18	Probabilistic Learning -I	Why generative models? Generative Vs Discriminative models, Bayes Rule, Maximum A Posteriori (MAP) Vs Maximum Likelihood (ML), Bayesian Networks: Building BNs, Inferencing, Direct, Serial, Converging and Diverging Connections, D-separations.	T1 (8.1, 8.4.1), T2(6)
19 - 21	Probabilistic Learning-II	Bayes to Naive Bayes: Why? Naive Bayes learning: Independence assumption, Types of Naive Bayes algorithms.	T1(4.2, 4.3), T2(6.1- 6.10)
22 - 24	Neural Networks-I	Perceptron learning algorithm, Multi-layer Network (MLP): Components, Activations, Training: Computing Gradients, Error Back-propagation.	T1(5.1- 5.4)

25 - 28	Neural Networks-II	Regularization, Data Augmentation, Convolutional Neural Networks: CNNs, RNNs; Generative AI Models: Autoregressive, GANs, VAEs.	T1(5.5), Lecture Slides
29 – 31	Instance-based and Kernel based Learning	k-Nearest Neighbour Learning, Constructing Kernels, Radial Basis Function Networks, Maximum margin classifiers (SVMs).	T2(8.2), T1(6.1-6.3, 7.1)
32 - 35	Un-supervised Learning - I	K-means Clustering Algorithm: Expectation Maximization, Convergence, Application of K-means, Gaussian Mixture Models: EM for GMM.	T1 (9.1, 9.2)
36 - 38	Un-supervised Learning - II	Dimensionality Reduction: What and Why? Principal Component Analysis (PCA) for feature reduction: Maths behind PCA.	T1(12.1)
39 - 41	Reinforcement Learning	Markov Decision Process, Value Iteration, Policy Iteration, Q- learning.	T1 (13.1), T2(13.3)

Evaluation Scheme:

Component	Duration	Date & Time	Weightage	Nature
				of
				Component
Mid-Semester	90 mins	08/10/2024 (9.30am-11.00am)	25%	Closed Book
Exam				
Home	5 to 6 nos.	1 to 2 assignments every month	30%	Open Book
Assignments/		for 4 months (Aug to Nov).		
Projects (coding)				
Two announced	30 mins	13/09 (6:00pm to 6:30pm)	10%	Open Book
quizzes	each	14/11 (6:00pm to 6:30pm)		
Comprehensive	3 Hrs	11/12/2024 (FN)	35%	Closed Book
Exam				

(Note: Minimum 40% of the evaluation component will be conducted before the mid semester grading)

Chamber Consultation Hours: Would be announced in the class.

Make-up Policy: Prior permission of the Instructor-In-Charge is required to get make-up on any evaluation component. Genuine requests will only be considered.

Notices: All notices about the course will be put on Course webpage.

Academic Honesty and Integrity Policy: Academic honesty and integrity are to be maintained by all the students throughout the semester and no type of academic dishonesty is acceptable.

Instructor-In-Charge, BITS F464