Proposed Course Syllabus

Title	**Algorithms for Big Data**	Number	CSL7030
Department	Computer Science and Engineering	L-T-P [C]	2–0–0 [2]
Offered for	M.Tech. 1st Year, Ph.D. 1st Year	Type	Compulsory
Prerequisite	None		

Objectives

The Instructor will:

1. Introduce some algorithmic techniques developed for handling large amounts of data.
2. Emphasize both theoretical as well as practical aspects of such algorithms.

Learning Outcomes

The students are expected to have the ability to:

1. Analyze existing algorithms as well as design novel algorithms pertaining to big data.

Contents

Introduction: Randomized algorithms, Universal Hash Family, Probabilistic Algorithm Analysis, Approximation Algorithms, ε-Approximation Schemes. (5 lectures)
Sketching and Streaming: Extremely Small-Space Data Structures, CountMin Sketch, Count Sketch, Turnstile Streaming, AMS Sketch, Graph Sketching, Graph Connectivity (9 lectures)
MapReduce: MapReduce Algorithms in Constrains Settings such as small memory, few machines, few rounds, and small total work, Efficient Parallel Algorithms (7 lectures)
External memory and cache-obliviousness: Minimizing I/O for large datasets, Algorithms and data structures such as B-trees, Buffer trees, Multiway merge sort (7 lectures)

Self Learning Material

1. Department of Computer Science, Harvard University, Algorithms for Big Data
2. https://www.sketchingbigdata.org/

Title	**Artificial Intelligence - I**	Number	CSL7XX0
Department	Computer Science and Engineering	L-T-P [C]	3–0–0 [3]
Offered for	M.Tech. 1st Year, Ph.D. 1st Year	Type	Compulsory
Prerequisite	None		

Objectives

The Instructor will:

1. Cover various paradigms that come under the broad umbrella of AI.

Learning Outcomes

The students are expected to have the ability to:

1. Develop an understanding of where and how AI can be used.

Contents

Introduction: Uninformed search strategies**,** Greedy best-first search, And-Or search, Uniform cost search, A* search, Memory-bounded heuristic search, Local and evolutionary searches (9 Lectures)
Constraint Satisfaction Problems: Backtracking search for CSPs, Local search for CSPs (3 Lectures)
Adversarial Search: Optimal Decision in Games, The minimax algorithm, Alpha-Beta pruning, Expectimax search (4 Lectures)
Knowledge and Reasoning: Propositional Logic, Reasoning Patterns in propositional logic; First order logic: syntax, semantics, Inference in First order logic, unification and lifting, backward chaining, resolution (9 Lectures)
Representation: Information extraction, representation techniques, foundations of Ontology (3 Lectures)
Planning: Situation Calculus, Deductive planning, STRIPES, sub-goal, Partial order planner (4 Lectures)
Bayesian Network, Causality, and Uncertain Reasoning: Probabilistic models, directed and undirected models, inferencing, causality, *Introduction to Probabilistic reasoning* (6 lectures)
Introduction to RL: MDP, Policy, Q-value (4 Lectures)

Textbook

1. *Russel,S., and Norvig,P., (2015), Artificial Intelligence: A Modern Approach, 3rd Edition, Prentice Hall*

Self Learning Material

1. Department of Computer Science, University of California, Berkeley, http://www.youtube.com/playlist?list=PLD52D2B739E4D1C5F
2. NPTEL: Artificial Intelligence, https://nptel.ac.in/courses/106105077/

Title	**Artificial Intelligence - II**	Number	CSL7XX0
Department	Computer Science and Engineering	L-T-P [C]	3–0–0 [3]
Offered for	M.Tech. 1st Year, Ph.D. 1st Year	Type	Compulsory
Prerequisite	None		

Objectives

1. To cover modern paradigms of AI that go beyond traditional learning

Learning Outcomes

The students are expected to have the ability to:

1. Develop an understanding of modern concepts in AI and where they can be used
2. Design, implement and apply novel AI techniques based on emerging real-world requirements

Contents

Making decisions: Utility theory, utility functions, decision networks, sequential decision problems, Partially Observable MDPs, Game Theory (14 Lectures)
Reinforcement Learning: Passive RL, Active RL, Generalization in RL, Policy Search, (7 Lectures)
Probabilistic Reasoning over time: Hidden Markov Models, Kalman Filters (7 Lectures)
Knowledge Representation: Ontological engineering, Situation Calculus, semantic networks, description logic (6 Lectures)
Planning: Planning with state space search, Partial-Order Planning, Planning Graphs, Planning with Propositional Logic, hierarchical task network planning, non-deterministic domains, conditional planning, continuous planning, multi-agent planning (8 Lectures)

Text Book

1. S. RUSSEL, P. NORVIG (2009), Artificial Intelligence: A Modern Approach, Pearson, 3rd Edition.

Reference Book

1. E. RICH, K. KNIGHT, S. B. NAIR (2017), Artificial Intelligence, McGraw Hill Education, 3rd Edition.
2. R.S. SUTTON, A.G. BARTO (2015), Reinforcement Learning: An Introduction, The MIT Press, 2nd Edition.

Title	**Machine Learning - I**	Number	CSL7XX0
Department	Computer Science and Engineering	L-T-P [C]	3–0–0 [3]
Offered for	M.Tech. (CSE, AI, DCS)	Type	Compulsory
Prerequisite	Introduction to Computer Sc., Probability, Statistics and Stochastic Processes	Antirequisite	IML, Applied ML, PRML

Objectives

1. To understand various key paradigms for machine learning approaches
2. To familiarize with the mathematical and statistical techniques used in machine learning.
3. To understand and differentiate among various machine learning techniques.

Learning Outcomes

The students are expected to have the ability to:

1. To formulate a machine learning problem
2. Select an appropriate pattern analysis tool for analyzing data in a given feature space.
3. Apply pattern recognition and machine learning techniques such as classification and feature selection to practical applications and detect patterns in the data.

Contents

Introduction: Definitions, Datasets for Machine Learning, Different Paradigms of Machine Learning, Data Normalization, Hypothesis Evaluation, VC-Dimensions and Distribution, Bias-Variance Tradeoff, Regression (Linear) (7 Lectures)
Bayes Decision Theory: Bayes decision rule, Minimum error rate classification, Normal density and discriminant functions (5 Lectures)
Parameter Estimation: Maximum Likelihood and Bayesian Parameter Estimation (3 Lectures)
Discriminative Methods: Distance-based methods, Linear Discriminant Functions, Decision Tree, Random Decision Forest and Boosting (5 Lectures)
Feature Selection and Dimensionality Reduction: PCA, LDA, ICA, SFFS, SBFS (4 Lectures)
Clustering: k-means clustering, Gaussian Mixture Modeling, EM-algorithm (4 Lectures)
Kernel Machines: Kernel Tricks, SVMs (primal and dual forms), K-SVR, K-PCA (6 Lectures)
Artificial Neural Networks: MLP, Backprop, and RBF-Net (4 Lectures)
Foundations of Deep Learning: DNN, CNN, Autoencoders (4 lectures)

Text Book

1. Shalev-Shwartz,S., Ben-David,S., (2014), Understanding Machine Learning: From Theory to Algorithms, Cambridge University Press
2. R. O. Duda, P. E. Hart, D. G. Stork (2000), Pattern Classification, Wiley-Blackwell, 2nd Edition.

Reference Book

1. Mitchell Tom (1997). Machine Learning, Tata McGraw-Hill
2. C. M. BISHOP (2006), Pattern Recognition and Machine Learning, Springer-Verlag New York, 1st Edition.

Self-Learning Material

1. Department of Computer Science, Stanford University, https://see.stanford.edu/Course/CS229

Title	**Machine Learning - II**	Number	CSL7XX0
Department	Computer Science and Engineering	L-T-P [C]	3–0–0 [3]
Offered for	M.Tech. (AI, DCS)	Type	Compulsory
Prerequisite	None	Antirequisite	Deep Learning

Objectives

1. Provide technical details about various recent algorithms and software platforms related to Machine Learning with specific focus on Deep Learning.

Learning Outcomes

The students are expected to have the ability to:

1. Design and program efficient algorithms related to recent machine learning techniques, train models, conduct experiments, and develop real-world ML-based applications and products

Contents

Fractal 1: Foundations of Deep Learning
Deep Networks: CNN, RNN, LSTM, Attention layers, Applications (8 lectures)
Techniques to improve deep networks: DNN Optimization, Regularization, AutoML (6 lectures)

Fractal 2: Representation Learning
Representation Learning: Unsupervised pre-training, transfer learning, and domain adaptation, distributed representation, discovering underlying causes (8 lectures)
Auto-DL: Neural architecture search, network compression, graph neural networks (6 lectures)

Fractal 3: Generative Models
Probabilistic Generative Models: DBN, RBM (3 lectures)
Deep Generative Models: Encoder-Decoder, Variational Autoencoder, Generative Adversarial Network (GAN), Deep Convolutional GAN, Variants and Applications of GANs (11 lectures)

Text Book

1. Goodfellow,I., Bengio.,Y., and Courville,A., (2016), Deep Learning, The MIT Press .

Reference Book

1. Charniak, E. (2019), Introduction to deep learning, The MIT Press.
2. Research literature.

Self Learning Material

1. https://www.deeplearningbook.org/

Current Syllabus

Algorithms for Big Data

Sketching and Streaming: Extremely small-space data structures (4 lectures)
Numerical linear algebra: Algorithms for big matrices, Regression, Low-rank approximation, Matrix completion (8 lectures)
Compressed Sensing: Sparse signals, Linear measurements, Signal recovery (8 lectures)
External memory and cache-obliviousness: Minimizing I/O for large datasets, Algorithms and data structures such as B-trees, Buffer trees, Multiway mergesort (8 lectures)

Machine Learning I:

Supervised Learning 1-0-0 [1]
Introduction: Motivation, Different types of learning, Linear regression, Logistic regression (2 lectures)
Gradient Descent: Introduction, Stochastic Gradient Descent, Subgradients, Stochastic Gradient Descent for risk minimization (2 lectures)
Support Vector Machines: Hard SVM, Soft SVM, Optimality conditions, Duality, Kernel trick, Implementing Soft SVM with Kernels (4 lectures)
Decision Trees: Decision Tree algorithms, Random forests (2 lectures)
Neural Networks: Feedforward neural networks, Expressive power of neural networks, SGD and Backpropagation (3 lectures)
Model selection and validation: Validation for model selection, k-fold cross-validation, TrainingValidation-Testing split, Regularized loss minimization (1 lectures)

Unsupervised Learning and Generative Models 1-0-0 [1]
Nearest Neighbour: k-nearest neighbour, Curse of dimensionality (1 lecture)
Clustering: Linkage-based clustering algorithms, k-means algorithm, Spectral clustering (3 lectures)
Dimensionality reduction: Principal Component Analysis, Random projections, Compressed sensing (2 lectures)
Generative Models: Maximum likelihood estimator, Naive Bayes, Linear Discriminant Analysis, Latent variables and Expectation-maximization algorithm, Bayesian learning (5 lectures)
Feature Selection and Generation: Feature selection, Feature transformations, Feature learning (3 lectures)

Computational Learning Theory and Deep Neural Networks 1-0-0 [1]
Statistical Learning Framework: PAC learning, Agnostic PAC learning, Bias-complexity tradeoff, No free lunch theorem, VC dimension, Structural risk minimization, Adaboost (7 lectures)
Foundations of Deep Learning: DNN, CNN, RNN, Autoencoders (7 lectures)

Machine Learning II:

Introduction to Deep Learning 1-0-0[1] Model Search: Optimization, Regularization, AutoML (4 lectures) Deep Networks: Attention layers, Gated CNNs, Graph Neural Networks (8 lectures) Applications: Neural language models (2 lectures)

Machine Learning II: Representation Learning & Structured Models 1-0- 0[1] Representation Learning: Unsupervised pre-training, transfer learning and domain adaptation, distributed representation, discovering underlying causes (7 lectures) Structured models: learning about dependencies, inference and approximate inference, sampling and Monte Carlo Methods, Importance Sampling, Gibbs Sampling, Partition Function, MAP inference and Sparse Coding, Variational Inference (7 lectures)

Machine Learning II: Deep Generative Models 1-0-0[1] Deep Generative Models: Deep Belief Networks, Variational Autoencoder, Generative Adversarial Network (GAN), Deep Convolutional GAN, Autoencoder GANs, iGAN, pix2pix, CycleGAN, Conditional GANs, StackGAN (14 lectures)

Artificial Intelligence - I

Probabilistic Reasoning over time: Hidden Markov Models, Kalman Filters, Dynamic Bayesian Networks (7 lectures)
Knowledge Representation: Ontological engineering, Semantic Networks, Description Logics (7 lectures)
Making decisions: Utility theory, utility functions, decision networks, sequential decision problems, Partially Observable MDPs, Game Theory (14 lectures)
Reinforcement Learning: Passive RL, Active RL, Generalization in RL, Policy Search, Deep Reinforcement Learning (14 lectures)

Artificial Intelligence - II

Introduction (1 lecture)
Propositional logic (8 lectures)
Search: Uninformed strategies (BFS, DFS, Dijkstra), Informed strategies (A* search, heuristic functions, hill-climbing), Adversarial search (Minimax algorithm, Alpha-beta pruning) (10 lectures)
Predicate logic: Knowledge representation, Resolution (6 lectures)
Rule-based systems: Natural language parsing, Context free grammar (3 lectures)
Constraint satisfaction problems (4 lectures)
Planning: State space search, Planning Graphs, Partial order planning (4 lectures)
Uncertain Reasoning: Probabilistic reasoning, Bayesian Networks, Dempster-Shafer theory, Fuzzy logic (6 lectures)