संगणक विज्ञान एवं अभियांत्रिकी विभाग
Department of Computer Science & Engineering
भारतीय प्रौद्योगिकी संस्थान जोधपुर
Indian Institute of Technology Jodhpur
एन. एच. 62, नागौर रोड़, करवर, जोधपुर – 342030
N.H. 62, Nagaur Road, Karwar, Jodhpur-342030
दूरभाष संख्या: +91 291 2801252 ईमेल: office_cse@iitj.ac.in

PhD CS Curriculum

PhD curriculum for students with B.Tech., B.S., B.E. or B.Sc. (Engineering), M.Sc. degree

Cat.	Course Number, Course Title		L-T-P	Credits	Cat.	Course Number, Course Title		L-T-P	Credits
I Semester					**II Semester**				
C	CSL7130	Mathematical Foundations for Computer Science	3-0-0	3	C	CSL7xx0	Compulsory*	3-0-0	3
C	CSL7xx0	Compulsory*	3-0-0	3	E	XXX7XX0	Elective**	3-0-0	3
C	CSL7xx0	Compulsory*	3-0-0	3	E	XXX7XX0	Elective**	3-0-0	3
E	XXX7XX0	Elective**	3-0-0	3	E	XXX7XX0	Elective**	3-0-0	3
NG1	HSN7XX0	Non-Graded I	1-0-0	1	NG2	HSN7XX0	Non-Graded II	1-0-0	1
Total				**13**	**Total**				**13**
III Semester					**IV Semester**				
T	CST8XX0	Thesis		12	T	CST8XX0	Thesis		12
Total				**12**	**Total**				**12**

S.No.	Category	Course Category Title	Total Courses	Total Credits
1	C	Compulsory	4	12
2	E	Electives	4	12
3	NG	Non-Graded	2	2
3	T	Thesis		
		Total		

PhD curriculum for students with M.Tech., M.E. or M.Sc. (Engineering)

Cat.	Course Number, Course Title		L-T-P	Credits	Cat.	Course Number, Course Title		L-T-P	Credits
I Semester					**II Semester**				
C	CSL7130	Mathematical Foundations for Computer Science	3-0-0	3	T	CST8XX0	Thesis		12
C	CSL7xx0	Compulsory*	3-0-0	3	NG2	HSN7XX0	Non-Graded II	1-0-0	1
C	CSL7xx0	Compulsory*	3-0-0	3					
C	CSL7xx0	Compulsory*	3-0-0	3					
NG1	HSN7XX0	Non-Graded I	1-0-0	1					
Total				**13**	Total				**13**
III Semester					**IV Semester**				
T	CST8XX0	Thesis		12	T	CST8XX0	Thesis		12
Total				**12**	Total				**12**

S.No.	Category	Course Category Title	Total Courses	Total Credits
1	C	Compulsory	4	12
2	E	Electives	0	0
3	NG	Non-Graded	2	2
3	T	Thesis		
		Total		

*All PhD students are required to take “Mathematical Foundations for Computer Science” as a compulsory course. In addition, they must choose any three courses from the Compulsory list as their other compulsory courses.

**PhD students can take electives from the Compulsory list (the courses that they have not taken as compulsory courses) or Electives list.

Compulsory List

S.No.	Course Number	Course Title	L-T-P	Credits
1	CSL7560	Advanced Data Structures and Algorithms	3-0-0	3
2	CSL7190	Advanced Architecture and OS	3-0-0	3
3	CSL7090	Software and Data Engineering	3-0-0	3
4	CSL7610	Artificial Intelligence	3-0-0	3
5	CSL7620	Machine Learning	3-0-0	3
6	CSL7590	Deep Learning	3-0-0	3

Electives List

S.No.	Course Number	Course Title	L-T-P	Credits
1	CSL7150	Cybersecurity	3-0-0	3
2	CSL7140	Complexity Theory	3-0-0	3
3	CSL7160	Advanced Algorithms	3-0-0	3
4	CSL7XXX	Emerging Network Technologies	3-0-0	3
5	CSL7XXX	Blockchain Technologies	3-0-0	3
6	CSL7360	Computer Vision	3-0-0	3
7	CSL7750	Distributed Database Systems	3-0-0	3
8	CSL7510	Virtualisation and Cloud Computing	3-0-0	3
10	CSL7460	Mobile and Pervasive Computing	3-0-0	3
11	CSL7640	Natural Language Understanding	3-0-0	3
12	CSL7XX0	High Performance Computing Architectures	3-0-0	3
13	CSL7650	Autonomous Systems	3-0-0	3
14	CSL7770	Speech Understanding	3-0-0	3
15	CSL7860	Foundation Models and Generative AI	3-0-0	3
16	CSL7870	Learning on Graphs and its Applications	3-0-0	3
17	CSL7110	Machine Learning with Big Data	3-0-0	3
18	CSL7210	ML Ops and DL-Ops	2-0-2	3
19	CSL7XXX	Selected Topics in AI	3-0-0	3
20	CSL7450	Computer Graphics	3-0-0	3
21	CSL7XXX	Systems for ML	3-0-0	3

Title	**Advanced Data Structures & Algorithms**	Number	**CS7560**
Department	Computer Science	L-T-P [C]	3–0–0 [3]
Offered for	M.Tech. CSE and M.Tech. AI	Type	core
Prerequisite	Algorithm Design and Analysis, Data Structures, Maths for Computing		

Objectives

1. The objective of the course is to introduce several advanced algorithmic techniques and data structures.
2. Implementation of various advanced data structure and algorithms to solve real life problems

Learning Outcomes

1. Learn advanced data structure techniques
2. Learn a new set of techniques to cope with NP-hard problems.
3. Identify novel and significant open research questions in the field.
4. Implementation of various advanced data structure and algorithms to solve real life problems
5. Understand the intractability of problems.

Contents

Data Structures and Algorithms Techniques (Fractals I and II)

Advanced Data Structures [10 lectures]: Red-black trees, Interval trees, Dynamic order statistics, Binomial heap, Fibonacci Heap [5 lectures], Tries, Splay Trees, B-Trees, Hashing - disjoint set, union find [5 lectures].
Network Flows [5 lectures]: Maximum Flow problem, Ford-Fulkerson algorithm, Max flow min cut theorem, augmenting paths, applications to bipartite matching problem and disjoint paths in directed and undirected graphs.
Amortized Analysis [3 lectures]: Aggregate Analysis, Accounting method, Potential method.
Dynamic Programming [6 lectures]: Knapsack, Independent Set of Trees, Interval Scheduling, Set Cover, Permutation Problems, Partition Problems, String matching: naive method, Knuth-Morris-Pratt algorithm.
NP-completeness and Reductions [8 lectures]: Complexity classes P, NP, co-NP, NP-complete, NP-Hard, Example reductions between problems.

Fractal III: Advanced Algorithmic Techniques

Advanced Algorithmic Techniques [10 lectures]: Approximation Algorithms: Greedy Algorithm – Load Balancing, Set Cover [2 Lectures]; The Pricing Method: Vertex Cover, Knapsack [3 Lectures], Randomized Algorithms [5 Lectures].

Textbooks

1. Jon Kleinberg, Eva Tardos (2005), Algorithm Design, Pearson Education, 1st Edition.
2. Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, Clifford Stein, Introduction to Algorithms.

Self Learning Material

1. https://ocw.mit.edu/courses/electrical-engineering-and-computer-science/6-046j-introduction-to-algorithms-sma-5503-fall-2005/video-lectures/
2. https://nptel.ac.in/courses/106/101/106101060/
3. https://nptel.ac.in/courses/106/104/106104019/

Title	**Advanced Architecture and Operating Systems**	Number	CSL7190
Department	CSE	L-T-P[C]	3–0–0 [3]
Offered for	M.Tech. CSE and M.Tech. AI	Type	Elective
Prerequisite	UG level OS and Computer Organization or Architecture	Antirequisite	None

Objectives

1. To develop foundations in advanced aspects of modern operating systems and architectures.
2. To develop capabilities in design thinking and systems design centered around the core layers of the computer.
3. To develop knowledge on end layer application use cases, performance, and usability.

Contents

Architecture: [24 lectures]

Processor Design [11 lectures]: Instruction Set Architecture (ISA) design, Advanced Pipelining, ILP, Out of order pipelines, Branch Prediction, Static and Dynamic Scheduling, Hardware Multithreading, Multicore Processors, Accelerators (ASIC/FPGA/GPU platforms).

Memory System: Memory Hierarchy [11 lectures], Advanced Cache Design, Instruction and data prefetching, Cache Coherence Protocols, Memory Consistency, Synchronization, non-uniform caches, Main memory design, Emerging Memory Technologies.

On-Chip Networks [2 lectures]: Message transmission, routing, performance.

Operating System: [15 lectures]

OS Fundamentals review, Types of OS, Types of Kernels, Design Approach, Virtualization (hardware-assisted), Containers and VM, Virtual machines scheduling,

Advanced Memory Management, Handling large pages, TLB-Memory Hierarchy Interactions, NUMA, Advanced File Systems, Distributed Operating Systems

Concurrency and Transactions, Security.

Textbooks

1. Arpaci-Dusseau, Remzi H., C. Andrea, and William Stallings. "Operating Systems: Three Easy Pieces."
2. Mukesh Singhal and Niranjan G. Shivaratri, “Advanced Concepts in Operating Systems Distributed, Database, and Multiprocessor Operating Systems”, Tata McGraw-Hill.
3. Hennessy, John L., and David A. Patterson. Computer architecture: a quantitative approach. Elsevier.

Reference

1. Research papers and reading material provided by instructors.

Title	**Software and Data Engineering**	Number	CSL7090
Department	Computer Science and Engineering	L-T-P [C]	3–0–0 [3]
Offered for	M.Tech., PhD	Type	Compulsory
Prerequisite			

Objectives

The Instructor will:

1. Discuss techniques to manage a large amount of data
2. Provide mechanisms to design and develop data-intensive computing systems

Learning Outcomes

The students are expected to have the ability to:

1. Design complex end-to-end data pipeline for data processing
2. Critically identify and use the tools for data handling and management
3. Use modern software technologies to design and develop data analytical systems

Contents

CSL7XX0 Cloud Computing and Virtualization 1-0-0 [1]

Basics of complex software design: Concept of modular software, microservices, communication, 4+1 architectural views and patterns (5 lectures)

Cloud Computing: Architecture of cluster computing, design of data centers, open data center platforms, fault-tolerant system design (5 lectures)

Virtualization: Type-1 and Type-2 virtualization, virtual machine, containers, dockers (4 lectures)

CSL7XX0 Data Management 1-0-0 [1]

Data Management: Structured data, relational database management, unstructured data, semi-structured data, Nosql database management (mongodb), column database, graph database, XML, JSON, HDFS, Handling drift in data, sensor data reliability at software and algorithmic level, sensor data analysis techniques (14 lectures)

CSL7XX0 Data Intensive Processing Systems 1-0-0 [1]

Data Intensive Processing Systems: Architecture of large scale data processing systems, Hadoop, Apache Spark, Storm, parallel data processing concepts such as map-reduce, directed acyclic graph, resilient distributed datasets, dynamic resource allocation, partial & shared computation, storage architecture (14 lectures)

Textbook

1. Bass L., Clements P., Kazman R., (2012), *Software Architecture in Practice*, 3rd edition, Addison-Wesley Professional
2. Martin K., (2017), *Designing Data-Intensive Applications: The Big Ideas Behind Reliable, Scalable, and Maintainable Systems*, 1st Edition, O'Reilly Media

Self Learning Material

1. Tylor,R.N., Medvidovic,N. and Dashofy,E.M., (2014), *Software Architecture Foundation: Theory and Practice*, Wiley

Preparatory Course Material

1. IEEE Transactions on Knowledge and Data Engineering
2. International Conference on Data Engineering

Title	**Artificial Intelligence (700)**	Number	CSL7610
Department	Computer Science and Engineering	L-T-P [C]	3–0–0 [3]
Offered for	B. Tech., M.Tech., Ph.D.	Type	Compulsory (AI)
Prerequisite	Data Structures and Algorithms	Antirequisite	Artificial Intelligence (300) - CSL 3XX

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)
Planning: Situation Calculus, Deductive planning, STRIPES, sub-goal, Partial order planner (3 Lectures)
Bayesian Network, Causality, and Uncertain Reasoning: Probabilistic models, directed and undirected models, inferencing, causality, *Introduction to Probabilistic reasoning* (6 lectures)
Reinforcement Learning: MDP, Policy, Q-value, Passive RL, Active RL, Policy Search (8 Lectures)

Textbook

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

Reference Books

1. Research literature

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	**Machine Learning**	Number	CSL7620
Department	Computer Science and Engineering	L-T-P [C]	3–0–0 [3]
Offered for	M.Tech. (CSE, AI, DCS), PhD	Type	Compulsory
Prerequisite	Introduction to Computer Sc., Probability, Statistics and Stochastic Processes, Linear Algebra	Antirequisite	IML, 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

Fractal I: Supervised Learning

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 (2 Lectures)

Fractal II: Unsupervised Learning

Discriminative Methods: Distance-based methods, Linear Discriminant Functions, Decision Tree, Random Decision Forest and Boosting (6 Lectures)

Feature Selection and Dimensionality Reduction: PCA, LDA, ICA, SFFS, SBFS (4 Lectures)

Clustering: k-means clustering, Gaussian Mixture Modeling, EM-algorithm (4 Lectures)

Fractal III: Kernels and Neural Networks

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	**Deep Learning (700)**	Number	CSL7590
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 (AI)
Prerequisite	Machine Learning	Antirequisite	Deep Learning (400) - CSL4xx

Objectives

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

Learning Outcomes

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 DL-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/

Course Title	**Cybersecurity**	Course No.	**CSL7150**
Department	Computer Science and Engineering	Structure (L-T-P [C])	3-0-2 [4]
Offered for	B. Tech, MTech, Ph. D.	Type	Elective
Prerequisite	Computer Networks	Anti-req	Security and its applications

Objectives

1. The Instructor will provide the skills needed to protect networks, secure electronic assets, prevent attacks, ensure the privacy of your customers, and build secure infrastructure

Learning Outcomes

1. To protect data and respond to threats that occur over the Internet
2. Design and implement risk analysis, security policies, and damage assessment
3. An ability to apply security principles and practices to the environment, hardware, software, and human aspects of a system.

Contents:

Introduction to Cyber Security: *Internet governance – Challenges and constraints, Cyber threats. (2 Lectures)*
Cyber Security Vulnerabilities and Cyber Security Safeguards: *Cyber security, Vulnerabilities, safeguards, Access control, Authentication, Biometrics, Deception, Denial of Service Filters, Ethical hacking, Firewalls, Response, Scanning, Security policy, Threat management. (8 Lectures)*
Cryptography: *Shannon's Approach to Cryptography: Measures of security, Perfect secrecy, Definition of entropy, One-time pad , Symmetric Key Cryptography, Cryptographic Hash Functions, Authentication, Public Key Cryptosystems, Key Distribution and Key Agreement Protocols(4 Lectures)*
Securing Web Application, Services and Servers: *Basic security for HTTP applications and services, Basic security for SOAP services, Identity management and Web services, Authorization patterns, Security considerations, challenges. (4 Lectures)*
Network Security: *TCP/IP threats, The IPSEC protocol, The SSL and TLS protocols, Firewalls and Virtual Private Networks (VPNs), Electronic mail security, Worms, DDoS attacks, BGB and security considerations (12 lectures)*
Intrusion Detection and Prevention: *Intrusion detection and Prevention techniques, Anti-malware software, Security information management, Network session analysis, System integrity validation. (4 Lectures)*
Cyberspace and the Law: *Cyber security regulations, Roles of international law, State and Private Sector in cyberspace, Cyber security standards. The INDIAN cyberspace, National cyber security policy. (5 Lectures)*
Cyber Forensics: *Handling preliminary investigations, Controlling an investigation, Conducting disk-based analysis, Investigating information-hiding, Tracing internet access, Tracing memory in real-time. (4 Lectures)*

Laboratory

1. Design and implementation of a simple client/server model and running application using sockets and TCP/IP.
2. To make students aware of the insecurity of default passwords, printed passwords and passwords transmitted in plain text.
3. To teach students how to use SSH for secure file transfer or for accessing local computers using port forwarding technique.
4. Comparison between Telnet and SSH for Secure Connection
5. AVISPA Tool for the Automated Validation of Internet Security Protocols and Applications

Text Book

1. C.J. HOOFNAGLE (2016), Federal Trade Commission Privacy Law and Policy, Cambridge University Press, 2016.
2. Stallings, W. (2017). Cryptography and network security, 7/E. Pearson Education India

Reference Book

1. Douglas R. Stinson, Maura B. Paterson (2018). Cryptography: theory and practice. 4/E Chapman and Hall/CRC
2. P.W. SINGER, A. FRIEDMAN (2014), Cybersecurity: What Everyone Needs to Know, OUP, 1st Edition.

Supplementary Resource

Title	**Complexity Theory**	Course No.	**CSL7140**
Department	Computer Science and Engineering	Structure (L-T-P [C])	3–0–0 [3]
Offered for	B.Tech., M.Tech., Ph.D. (CSE, AI), M.Sc. (Maths)	Type	Core
Prerequisite	Design and Analysis of Algorithms, Maths for Computing (or Discrete Mathematics)	Antirequisite	

Objectives:

1. Learning the limits of computation.
2. Learning basic complexity classes and various computational models.

Learning Outcomes:

1. A Complexity Theory course covers computational complexity and how to analyze the time and space efficiency of algorithms using big-O notation. Students learn about complexity classes like P, NP, NP-hard, and NP-complete, and how to establish problem hardness using reduction techniques.
2. The course introduces complexity hierarchies, alternation, probabilistic complexity. Students gain skills in classifying problems and understanding the limits of efficient computation.

Contents

Computability [7 lectures]: Review of Turing Machine (TM) and its variants, Universal Turing Machines and encoding TMs, Uncomputability, Complexity Class P.

Time Complexity [7 lectures] : Time as a resource, the Complexity Class NP, P vs NP, NP-completeness, Cook-Levin Theorem, Reductions.

Diagonalization [3 lectures]: Time-hierarchy theorem, Ladner's theorem, Mahaney's theorem.

Space Complexity [7 lectures]: Space as a resource, the Complexity Classes PSPACE, L and NL, Reachability Problem, PSPACE and NL completeness, Space Hierarchy Theorem, Savitch's theorem, Immerman-Szelepscenyi theorem.

Polynomial Hierarchy [4 lectures]: The Polynomial Hierarchy, Alternating TMs, Time-Space Tradeoffs for the Satisfiability problem.

Circuits [7 lectures]: Boolean Circuits and the Class P/poly**,** Advice taking TMs, Karp-Lipton Theorem, Circuit Lower Bounds, the Complexity Classes NC and AC.

Randomized Computation [7]: Probabilistic TMs, the Complexity Classes RP, coRP and ZPP, Error-reduction in randomized computation, Sipser-Gacs theorem, Randomized bounded-space computation.

Textbooks

1. S. Arora and B. Barak (2009), Complexity Theory: A Modern Approach, Cambridge University Press.

Reference books

1. Sipser, M., (2013), Introduction to the Theory of Computation, Cengage Learning.
2. Oded Goldreich (2008), Computational Complexity: A Conceptual Perspective.

Self Learning Material

1. https://www.youtube.com/playlist?list=PLm3J0oaFux3YL5vLXpzOyJiLtqLp6dCW2
2. https://onlinecourses.nptel.ac.in/noc21_cs49/preview

Title	**Advanced Algorithms**	Number	**CSL7160**
Department	Computer Science	L-T-P [C]	3–1–0 [4]
Offered for	B.Tech., M.Tech., Ph.D. (CSE, AI), M.Sc. (Maths)	Type	Elective / Core Specialization
Prerequisite	Design and Analysis of Algorithms, Maths for Computing (or Discrete Mathematics)		

Objectives

1. The objective of the course is to introduce several advanced algorithmic techniques such as parameterized algorithms, approximation algorithms, randomized algorithms, etc.

Learning Outcomes

1. Learn a new set of techniques to cope with NP-hard problems.
2. Identify novel and significant open research questions in the field.

Contents

Parameterized Algorithms [6 lectures]: Introduction to Parameterized Complexity and basics techniques: Branching, Iterative Compression.

Approximation Algorithms: [10 lectures]: Greedy Algorithm – Load Balancing, Center Selection Problem, Set

Cover [5 Lectures]; The Pricing Method: Vertex Cover, Linear Programming and Rounding: An application to Vertex Cover, Knapsack [5 Lectures]

Randomized Algorithms [10 lectures]: Contention Resolution, Global, Random Variables and Expectations, Max-3-SAT approximation [7 Lectures]; Color Coding [3 Lectures]

Exact Exponential Time Algorithms [7 lectures]: Exact Algorithms for Coloring, SAT, Directed Feedback Arc Set, Max-Cut, Monotone-Local-Search, Or some other topics of contemporary interest.

Streaming Algorithms [5 lectures]: Introduction to streaming algorithms and its application to some graph theoretic problems.

Online Algorithms [4 lectures]: Introduction to online algorithms: definition, the difference with offline setting; list update problem, paging problem, Bin packing.

Textbooks

1. Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Daniel Marx, Marcin Pilipczuk, Michal Pilipczuk, Saket Saurabh (2015): Parameterized Algorithms, Springer.
2. Jon Kleinberg, Eva Tardos (2005), Algorithm Design, Pearson Education, 1st Edition.
3. Fedor V. Fomin, Dieter Kratsch (2010), Exact Exponential Time Algorithms, An EATCS Series, Springer.

Self Learning Material

1. https://www.youtube.com/watch?v=Ex8TueBsF1g&list=PLhkiT_RYTEU0gpi97fqjtaHy9Gk47oF85&index=1
2. https://sites.google.com/view/sakethome/teaching/parameterized-complexity?authuser=0
3. https://www.youtube.com/watch?v=S8Acu3EpvsE&list=PLhkiT_RYTEU2itsMqCNdXUg4cdFUWJn3-&index=4
4. https://www.youtube.com/watch?v=jNfQ3GZlrjM&list=PLhkiT_RYTEU3vSaVleEm_-bIPBzCqRQHK
5. https://www.youtube.com/watch?v=0rB6Pyq2cLQ&list=PLTtM9ThZ2L-c638AjqTskivmXwVTzTYnl&index=1

Title	**Blockchain Technology**	Course No.	CSL 7XX
Department	Computer Science and Engineering	Structure (L-T-P [C])	3–0–0 [3]
Offered for	B.Tech, MTech, PhD.	Type	Elective
Prerequisite	Cryptography	Antirequisite	

Objectives

The Instructor will:

1. Explain how blockchain technology works
2. Integrate blockchain technology into the current business processes to make them secure

Learning Outcomes

1. Understand what and why of Blockchain
2. Explore major components of Blockchain and Identify a use case for a Blockchain application Create their own Blockchain network application

Contents

Introduction to Blockchain: Digital Trust, Asset, Transactions, Distributed Ledger Technology, Types of network, Components of blockchain, mining, double spending attack, 51% attack. (5 Lectures)
PKI and Cryptography: Private keys cryptography, Public keys cryptography, Hashing, Digital Signature. (6 Lectures)
Consensus: Byzantine Fault, Proof of Work, Proof of Stake. (5 Lectures)
Cryptocurrency: Bitcoin creation and economy, Limited Supply and Deflation, Hacks, Ethereum concept and Ethereum classic, Monero, Zcash, Zero knowledge proofs, etc. (9 Lectures)
Permissioned blockchains: Permissionless versus permissioned Blockchains (2 Lectures)
Hyperledger Fabric: Hyperledger Architecture, Membership, Blockchain, Transaction, Chaincode, Hyperledger Fabric, Features of Hyperledger, Fabric Demo. (8 Lectures)
Blockchain Applications: Building on the Blockchain, Ethereum Interaction - Smart Contract and Token (Fungible, non-fungible), Languages,Payment escrow, Micro payments, Decentralized lotteries Blockchain-as-a-service. (6 Lectures)

Textbook

A. BAHGA, V. MADISETTI (2017), Blockchain Applications: A Hands-On Approach, VPT.

Self Learning Material

1. M. SWAN (2015), Blockchain: Blueprint for a New Economy, O'Reilly Media.
2. R. WATTENHOFER (2016), The Science of the Blockchain, CreateSpace Independent Publishing Platform.
3. I. BASHIR (2017), Mastering blockchain, Packt Publishing Ltd.
4. K.E. LEVY, Book-smart, Not Street-smart: Blockchain-based Smart Contracts and the Social Workings of Law, Engaging Science, Technology, and Society, vol. 3, pp. 1-15, 2017.

Preparatory Course Material

MIT Online Blockchain Course, Learn Blockchain Technology: https://getsmarter.mit.edu/

Title	**Computer Vision**	Number	CSL7360
Department	Computer Science and Engineering	L-T-P [C]	3–0–0 [3]
Offered for	B.Tech, M.Tech., PhD	Type	Elective
Prerequisite			

Objectives

The Instructor will:

1. Provide insights into fundamental concepts and algorithms behind some of the remarkable success of Computer Vision
2. Impart working expertise by means of programming assignments and a project

Learning Outcomes

The students are expected to have the ability to:

1. Learn and appreciate the usage and implications of various Computer Vision techniques in real-world scenarios
2. Design and implement basic applications of Computer Vision

Contents

Introduction: The Three R's - Recognition, Reconstruction, Reorganization (1 lecture)
Perspective: Static Perspective, Transformations, Dynamic perspective (5 lectures)
Fundamentals of Image formation and processing: Radiometry of image formation, Basic image processing, Biological visual processing (5 lectures)
Recognition: Object recognition case study - identifying digits with multiple approaches, Visual grouping, Convolutional Neural Network (ConvNet) based approaches to visual recognition of objects and scenes, Deformable Parts Model (DPM), Attributes, pose and actions (8 lectures)
Analysis: Binocular Stereopsis, Markov Random Fields in Computer Vision, Solving for stereo correspondence, Optical flow, Review of differential geometry (12 lectures)
Detection and Segmentation: Contour detection, Bottom-up segmentation, Gestalt grouping heuristics, Semantic segmentation - instance segmentation and pixel classification, Pose and keypoint estimation (5 lectures)
Image understanding: Scene understanding from RGBD images, 3D perception from a single image, Face recognition (6 lectures)

Textbook

1. Hartley,R. Zisserman,A., (2004), *Multiple View Geometry in Computer Vision*, 2nd Edition, Cambridge University Press
2. Szeliski,R., (2010), *Computer Vision: Algorithms and Applications*, Springer-Verlag London

Reference Books

1. Research literature

Course Title	**Distributed Database Systems**	Course No.	**CSL7750**
Department	Computer Science & Engineering	Structure (L-T-P [C])	3-0-0 [3]
Offered for	B.Tech, MTech, PhD	Type	Elective
Prerequisites	Operating Systems, Database Systems, Computer Networks	Antirequisite	None

Objective
To understand and appreciate concepts of distributed database design, and its associated issues of consistency, concurrency, optimization, integrity, reliability, privacy, and security.

Learning Outcome
Ability to understand the need for distributed database systems and its related complexities pertaining to fragmentation, replication, availability, concurrency, consistency and recovery.

Contents
Introduction: Distributed data processing concepts, What is a DDBS - advantages, disadvantages and problem areas. (2 Lectures)
Distributed Database Management System Architectures: Transparencies, Architecture, Global directory concepts and issues. (3 Lectures)
Distributed Database Design: Design strategies, Design issues, Fragmentation, Data allocation. (4 Lectures)
Semantics Data Control: View management, Data security, Semantic integrity control. (5 Lectures)
Query Processing: Objectives, Characterization of processors, Layers of processing, Query decomposition, Data localization. (5 Lectures)
Query Optimization: Factors, Centralized query optimization, Fragmented query ordering, Query optimization algorithms. (5 Lectures)
Transaction Management: Goals, Properties, Models. (4 Lectures)
Concurrency Control: Concurrency control in centralized systems, Concurrency control in DDBSs - algorithms, Deadlock management. (5 Lectures)
Reliability: Issues and types of failures, Reliability techniques, Commit protocols, Recovery protocols. (5 Lectures)
Other Avenues: Parallel database Systems, Multi-databases. (4 Lectures)

Reference Books
S. CERI, G. PELAGATTI (2008), Distributed Databases: Principles and Systems, McGraw-Hill, 1st Edition (2017 Reprint).
M.T. ÖZSU, P. VALDURIEZ (2011), Principles of Distributed Database Systems, Springer, 3rd Edition.

Course Title	**Virtualization and Cloud Computing**	Course No.	**CSL7510**
Department	Computer Science & Engineering	Structure (L-T-P [C])	3-0-0 [3]
Offered for	B.Tech	Type	Elective
Prerequisite	Operating Systems, Computer Networks	Antirequisite	None

Objectives
This course will introduce fundamentals of virtualization and concepts, frameworks, and applications in Cloud Computing to the audience

Learning Outcomes
1. To understand various virtualization techniques and their features and limitations
2. To understand design and concepts of a cloud computing framework

Contents
Cloud Computing: Concept, Definition, Cloud Types and Service Deployment Models. (5 Lectures)
Virtualization: Concept, Definition, Types of Virtualization, Hardware Virtualization, Full and Para Virtualization, Hypervisors, Hardware-assisted virtualization, operating system level virtualization, application virtualization (10 Lectures)
Virtual and Physical Networking: Introduction, vSwitches, virtual NICs, Virtual Networking, virtual LAN (5 Lectures)
Storage Virtualization: Introduction, SAN/NAS versus storage virtualization. (4 Lectures)
Virtual Machine Management: Base Virtual Machine, Virtual CPUs, Sockets, Cores, Memory Scaling Up and Scaling Down, USB Support, Virtual Disks, Live Migration. Security. (10 Lectures)
Containers: Concept, Definition, Docker, Container versus Virtualization, Portability, Remote deployment (5 Lectures)
Applications and Case Studies: Linux KVM, VirtualBox, Openstack. (3 Lectures)

Text Book
1. D.E. SARNA (2010), Implementing and Developing Cloud Computing Applications, CRC Press.
2. B.S. SODHI (2017), Topics in Virtualization and Cloud Computing, Ropar PB India, 2017.

Reference Book
B. FURHT, A. ESCALANTE (2010), Handbook of Cloud Computing (Vol. 3), Springer.

Course Title	**Mobile and Pervasive Computing**	Course No.	**CSL7460**
Department	Computer Science and Engineering	Structure (L-T-P [C])	3-0-0 [3]
Offered for	B. Tech and M. Tech	Type	Elective
Pre-requisites	Computer Networks	Antirequisite	None

Objectives

1. Explain necessary concepts to understand challenges in developing applications for mobile devices, when contrasted with that of a general-purpose computer.
2. Introduce mechanisms employed by different cellular technologies, particularly 5G, and contrast them with each other.

Learning Outcomes

At the end of this course, students would be able to:

1. Describe the working principles of various cellular technologies, as well as appreciate the associated challenges.
2. Compare/contrast various aspects of Mobile and Pervasive Computing, viz. localization, sensing, security and privacy concerns, application development, etc.
3. Implement the theoretical concepts into practical scenarios..

Contents

Overview of Wireless Systems: Infrastructure-based vs Ad-hoc, Wireless LANs, Cellular systems, Sensor networks, Bluetooth, WiFi, WiMAX. (3 Lectures)
Medium Access: Link Adaptation, Routing Protocols. (4 Lectures)
Mobility and Handoff Management: Link layer mobility mechanisms (location management protocols), Network layer mobility mechanisms (Macro and Micro mobility protocols), Handoff management protocols. (6 Lectures)
Cellular Networks: LTE and 5G overview, 5G Architecture, RAN and dynamic CRAN, Mobility management and Network slicing in 5G. (11 Lectures)
Pervasive Computing: Principles, Characteristics, Pervasive devices, Smart sensors and actuators, Context communication and access services. (4 lectures)
Context Aware Sensor Networks: Open protocols (SDP, Jini, SLP, UPnP), SyncML framework, Context aware mobile services, Context aware security. (4 lectures)
Energy Efficiency, Localization* (GPS/WiFi/GSM), *Security (4 Lectures)
Recent Advances in Wearable devices (3 Lectures)
***Body Area Networks* (BAN)** (3 Lectures)

Textbooks

1. I. STOJMENOVIC (2002), Handbook of Wireless Networks and Mobile and Pervasive Computing, Wiley.
2. S. LOKE (2006), Context-aware Pervasive Systems: Architectures for a New Breed of Applications, CRC Press.
3. A. OSSEIRAN, J.F. MONSERRAT, P. MARSCH, (Eds.), 5G Mobile and Wireless Communications Technology, Cambridge University Press, 2016.

Reference Books

1. R. KAMAL (2008), Mobile and Pervasive Computing, Oxford University Press, 3rd Edition.
2. L. MERK, M. NICLOUS (2006), Principles of Mobile and Pervasive Computing, Dreamtech Press, 2nd Edition.
3. G. AGGELOU (2004), Mobile Ad hoc Networks: From Wireless LANs to 4G Networks, McGraw-Hill Professional.

Self-learning Material:

1. J.J. DRAKE, Z. LANIER, C. MULLINER, P.O. FORA, S.A. RIDLEY, G. WICHERSKI (2014), Android Hacker's Handbook, John Wiley, 1st Edition.
2. P. SINGH, NPTEL, IIIT Delhi: https://nptel.ac.in/courses/106106147/
3. Relevant research papers from MobiSys, Sensys, MC2R, MobiCom and UbiComp.

Title	**Natural Language Understanding**	Number	CSL7640
Department	Computer Science and Engineering	L-T-P [C]	3–0–0 [3]
Offered for	B.Tech, M.Tech., PhD	Type	Elective
Prerequisite			

Objectives

The Instructor will:

1. To provide insights into fundamental concepts and algorithms related to Natural Language Understanding
2. Impart working expertise by introducing practical problems.

Learning Outcomes

The students are expected to have the ability to:

1. Formulate natural language understanding tasks
2. Design and implement basic applications of NLU

Contents

Traditional NLU: Introduction to NLU, Motivation, Morphology, Parts-of-Speech, Language Models, Word Sense Disambiguation, Anaphora Resolution, Basics of Supervised and Semi-supervised Learning for NLU, Hidden Markov Models for language modeling, EM Algorithm, Structured Prediction, Dependency Parsing, Topic Models, Semantic Parsing, Sentiment analysis. (14 Lectures)

Deep Learning for NLU: Intro to Neural NLU, Word Vector representations, Neural Networks and backpropagation -- for named entity recognition, Practical tips: gradient checks, overfitting, regularization, activation functions, Recurrent neural networks -- for language modeling and other tasks, GRUs and LSTMs -- for machine translation, Recursive neural networks -- for parsing, Convolutional neural networks -- for sentence classification, Question answering and dialogue system, Graph Neural Network for NLU, Natural Language Generation, Analysis and Interpretability of Neural NLU. (22 Lectures)

Know ledge Graphs: Knowledge graph embedding techniques, Inference on knowledge graphs. (6 Lectures)

Textbook

1. C. MANNING, H. SCHÜTZE (1999), Foundations of Statistical Natural Language Processing, MIT Press.
2. D. JURAFSKY, J.H. MARTIN, Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics and Speech Recognition (3rd Edition Draft), 2019.

Reference Books:

1. E. BENDER (2013), Linguistic Fundamentals for NLP, Morgan Claypool Publishers..
2. J. ALLEN (1995), Natural Language Understanding, Pearson Education, 1995.
3. Research Literature.

Self-Learning Material

1. http://web.stanford.edu/class/cs224n/index.html#schedule (Deep learning for NLP)

Title	**Autonomous Systems**	Number	CSL7650
Department/IDRP	CSE	L-T-P [C]	3-0-0 [3]
Offered for	B. Tech, MTech, PhD (All Branches)	Type	Elective
Prerequisite	Basics of Linear Algebra, Probability and Statistics		

Objectives

1. Provide an understanding about autonomous/ semi-autonomous systems like autonomous mobile robots and cars.

Learning Outcomes

1. Understand and use the methodologies to design, model and implementation of autonomous systems for real time applications.

Contents:

Fractal I: Introduction to Autonomous Systems [14 Lectures]

- Introduction to Autonomous Systems, properties of autonomous systems, examples of autonomous systems, degree of autonomy **[3]**
- Abstract Agent Architectures, Environment and agent, reactive, utility-based and behavior based architectures. Concrete agent architectures, logical agents, introduction to Bayesian reasoning **[7]**
- Introduction to real-time systems, real-time environment, distributed solutions, notion of time, fault-tolerance, modeling real-time systems **[4]**

Fractal II: Design of Autonomous vehicles [14 Lectures]

- Architecture of autonomous vehicles, processing pipeline, Vehicle Networks **[1]**
- Sensors for Navigation, GPS, IMU, Lidar, Camera, Visual Odometry, Place Recognition, Extraction based on Range Data **[5]**
- Localization, Noise and Aliasing, Belief Representation, Map Representations, Probabilistic Map based localization, Autonomous Map Building, SLAM **[4]**
- Planning and Reacting, Path Planning and Obstacle Avoidance, Navigation Architectures, Data Fusion **[4]**

Fractal III: Multi-agent systems [14 Lectures]

- Multi-agent systems and organizations, Hierarchical systems & peer-to-peer organization, Openness & scalability, Emergent organization **[2]**
- Communication and coordination in distributed systems, Distributed Planning and Execution, Task sharing, Coordination without communication [3]
- Game Theory, 2-player games, Nash Equilibrium, Strategy, Adversarial search, Iterative game and learning, Pareto-efficiency, Negotiation, Social choice, Mechanism design, Auction **[7]**
- Autonomous systems and human society, ethics and trust in multi-agent systems, explainability **[2]**

Text Books

1. Gerhard Weiss ed. (2013), Multiagent System, Second Edition, MIT Press.
2. Roland Siegwart, Illah Reza Nourbakhsh, Davide Scaramuzza (2018), Introduction to autonomous mobile robots, MIT press.

Reference Books

1. Michael Wooldridge (2009), An introduction to multi-agent systems (2nd ed). Wiley
2. Stuart J. Russell and Peter Norvig (2019), Artificial Intelligence: A Modern Approach, 4th edition,Pearson Press
3. Hermann Kopetz. (2011). Real-time Systems (2nd ed). Kluwer Academic Publishers
4. Choset, H. M., Hutchinson, S., Lynch, K. M., Kantor, G., Burgard, W., Kavraki, L. E., &Thrun, S. (2005). Principles of robot motion: theory, algorithms, and implementation. MIT press.

Other study materials

Research papers: (to be shared in class)

Online Course Material:

https://www.edx.org/course/autonomous-mobile-robots

Course Title	**Speech Understanding**	Course No.	**CSL7770**
Department	CSE / EE	Structure (L-T-P-C)	3-0-0 [3]
Offered for	B.Tech CSE, AI&DS	Type	Elective
Prerequisite	PRML/IML/ML	Antirequisite	None

Objectives
1. To provide insights into fundamental concepts and algorithms related to speech processing and understanding
2. Impart working expertise by introducing practical problems.

Learning Outcomes
Students will have the ability to:
1. Build a speech recognition system
2. Design and implement basic speech based application

Contents
Introduction to Speech processing: Digitization and Recording of speech signal, Review of Digital Signal Processing Concepts, Human Speech production, Acoustic Phonetics and Articulatory Phonetics, Different categories speech sounds and Location of sounds in the acoustic waveform and spectrograms.(14 Lectures)
Speech recognition*:* Analysis and Synthesis of Pole-Zero Speech Models, Short-Time Fourier Transform, Analysis:- FT view and Filtering view, Synthesis:-Filter bank summation (FBS) Method and OLA Method, Features Extraction, Extraction of Fundamental frequency, Speech Enhancement, Clustering and Gaussian Mixture models, Speaker Recognition.(14 Lectures)
Speech based applications: HMM and Neural models for speech recognition, Speech generation, Question answering, Dialogue systems, Other Speech based Applications.(14 Lectures)

Text Books
1. T.F. QUATIERI (2002), Discrete-Time Speech Signal Processing, Prentice-Hall, New Jersey.
2. D. JURAFSKY, J.H. MARTIN, Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics and Speech Recognition (3rd Edition Draft), 2019.

Reference Books
1. Y. GOLDBERG (2016), A Primer on Neural Network Models for Natural Language Processing, Journal of Artificial Intelligence Research.
2. I. GOODFELLOW, Y. BENGIO, A. COURVILLE (2016), Deep Learning, The MIT Press, 1st Edition.
3. S.K. PATRA (2011), Digital Signal Processing: A Computer-Based Approach, McGraw-Hill, 4th Edition.

Title	**Foundation Models and Generative AI**	Number	CSL7860
Department	Computer Science and Engineering	L-T-P [C]	3-0-0 [3]
Offered for	B. Tech, MTech, PhD (All Branches)	Type	Elective
Prerequisite	IML/PRML/ML and Deep Learning	Antirequisite	None

Objectives

1. This course aims to equip students with a deep understanding of foundation models, focusing on its fundamental concepts and core models. Students will learn about the concept and impact of foundation models within the generative AI domain. The course will explore the extensive capabilities of pre-trained models and their applications in AI-powered solutions. Additionally, students will gain practical insights into the features, capabilities, and real-world applications of various generative AI platforms. Students will be able to critically assess and effectively implement generative AI technologies in diverse contexts.

Learning Outcomes

1. The students will be able to understand the principles and formulations of foundation models,
2. The students will be able to conceptualize and design new foundation models, and utilize them for different applications and tasks,
3. The students will be able to analyze foundation models from different perspectives such as performance, computational requirements, explainability and interpretability,
4. The students will be able to perform inference and fine-tune foundation models.

Contents

Introduction to Foundation Models:

- Transformers & Attention : Introduction to Transformers, Self-Attention, Cross-Attention, Perceiver: Arbitrary IO with transformers, Self Attention & Non-Parametric Transformers (6 lectures)
- Diffusion : Sampling, Diffusion Models, Latent Diffusion Models (LDM) (6 lectures)
- Applications : LLMs, LVMs and Applications in Vision (2 lectures)

Foundation Model Architectures:

- Foundation Model Architectures: Transformers in Language including GPT, BERT, RetNet, State Space Models (4 lectures)
- Transformers in Vision: ViT, MLP-Mixer, Conformer, Vision-Language(VL) models, CLIP (4 lectures)
- Architectures: Dual-Encoder, Fusion, Encoder-Decoder, Adapted LLM, Unified Architectures (6 lectures)

Multi-Modal & Generative Foundation Models:

- Training & Inference: Training Objectives, Contrastive Learning, Efficient Inference Techniques, Pre-training, Fine-tuning & Parameter Efficient Fine-tuning (LoRA, QLoRa), Flash Attention, Retrieval Augmented Generation (RAG) (8 lectures)

Evaluation & Benchmarking:

- Evaluation Protocols and Standard Benchmarks, Hallucinations, Bias & Fairness, Privacy, Memorization, Machine Unlearning, Prompt Engineering (6 lectures)

Evaluation Plan:

- 3 exams and quizzes = 15+15+15+5 = 5o%
- Coding Project with three continuous evaluation components: 1o+1o+2o = 4o%
- Paper Reviews and class presentation: 5+5 = 1o%

Textbook & Reference Material

1. Foundation Models for Natural Language Processing - Pre-trained Language Models Integrating Media, Gerhard Paaß , Sven Giesselbach, Springer 2o23.
2. Research papers will be shared in class

Course Title	**Learning on Graphs and its Applications**	Course No.	**CSL7870**
Department	Computer Science and Engineering	Structure (L-T-P [C])	3-0-0 [3]
Offered for	B.Tech, MTech, PhD	Type	Elective
Pre-requisites	IML/PRML for UG	Preferred Knowledge	Basic ML

Objectives

Complex data can be modeled as a graph of relationships and interactions between objects. Graph data structures can be processed by algorithms like neural networks to perform tasks such as classification, clustering, and regression.This course investigates the computational, algorithmic, and modeling challenges specific to graph analysis. By studying underlying graph structures, students will gain mastery of Graph Neural Network techniques that can enhance prediction and uncover insights across diverse networks.

Learning Outcomes

The students are expected to have the ability to:

1. Understand the applications related to graph and graph representation learning
2. Formulate real-world problems with relational data and use GNN for prediction, classification and other tasks on graph data structures

Contents

Introduction: Complex Systems, Graph Data Structure, Graph Representations, Spectral graph theory , Real world networks and applications. (6 Lectures)

Graph Representation Learning: Node Embeddings , Graph Convolutional Networks, Design Space, Inductive Representation Learning, (6 Lectures)

Graph Attention Networks, Hierarchical Graph Representation Learning, Graph Embeddings, Over smoothing Issues, Expressiveness of GNN, GNN vs CNN (8 Lectures)

GNN Applications: Heterogeneous and Multilayer graphs, Knowledge graphs, Reasoning over knowledge graphs, Graphs and LLMs (8 Lectures)

Neural subgraph matching, GCN for Recommender Systems, Node Classification, Graph Classification, (6 Lectures)

Community Detection, Link Prediction and Causality. (4 Lectures)

Large Scale Graphs: Dealing with large graphs, Cluster-GCN, Simplifying Graph Convolutional Networks (4 Lectures)

Textbook

1. Graph Representation Learning by William L. Hamilton
2. Networks, Crowds, and Markets: Reasoning About a Highly Connected World , by David Easley and Jon Kleinberg, (Cambridge University Press - Sep 2010) – Pre-publication draft available online.

Reference Books

1. Network Science , by Albert-Laszlo Barabasi, (Cambridge University Press - August 2016) – freely available under the Creative Commons license.

Online Materials:

1. http://web.stanford.edu/class/cs224w/
2. https://youtube.com/playlist?list=PLoROMvodv4rPLKxIpqhjhPgdQy7imNkDn&si=39XL6M_ZzZSSD10j

Title	**Machine Learning with Big Data**	Number	CSL7110
Department	Computer Science and Engineering	L-T-P [C]	3–0–0 [3]
Offered for	B.Tech, M.Tech., PhD	Type	Elective
Prerequisite	Machine Learning / PRML	Antirequisite	None

Objectives

The Instructor will:

1. Provide an understanding of the role of big data in the real-world scenarios
2. Provide technical details about various algorithms and software/hardware tools/platforms related to big data

Learning Outcomes

Students are expected to have the ability to:

1. Develop an understanding of big data in the modern context, and independently work on problems relating to big-data
2. Design and program efficient algorithms for big data from the perspective of a project

Contents

Introduction: What is big data, Unreasonable effectiveness of data (1 lecture)
Streaming algorithms: Streaming Naive Bayes, Stream and sort (2 lectures)
Platforms for learning from big data: MapReduce, New Software Stack, Large Scale File System Organization (5 lectures)
Nearest Neighbour Search, Jaccardi Similarity of Sets, Similarity of Documents, Locality Sensitive Hashing, The Stream Data Model (4 lectures)
Randomized methods: Clustering, Hashing, Sketching, Scalable stochastic gradient descent (3 lectures)
Frequent Itemsets: The Market Basket Model, A-Priori Algorithm, Handling larger datasets in Main Memory, Limited-Pass Algorithms, Counting Frequent Items in a Stream (6 lectures)
Parameter Servers: Introduction, Abstraction, Parameter Cache Synchronization, Asynchronous execution, Model Parallel Examples (3 lectures)
Graph-based methods (6 lectures)
Page Rank, Topic Sensitive Page Rank, Approaches to Page Rank iteration, Link Spam, Semi-supervised learning, Scalable link analysis, Models for Recommendation Systems, Social Networks as Graphs (9 lectures)
Large-scale Machine Learning with CPUs and GPUs (3 lectures)

Textbook

1. Leskovec,J., Rajaraman,A., Ullman,J., (2014), *Mining of Massive Datasets*, 2nd Edition, Cambridge University Press
2. Bekkerman,R., Bilenko,M., Langford,J., (2011), *Scaling Up Machine Learning*, Cambridge University Press

Reference Books

1. Research literature

Self Learning Material

1. Department of Machine Learning, Carnegie Mellon University, Machine Learning with Large Datasets Course
2. Department of Computer Science, University of California, Berkeley, Scalable Machine Learning
3. ETH Zurich, Data Mining: Machine Learning from Large Datasets

Title	**ML-DL-Ops**	Number	**CSL7210**
Department	Computer Science	L-T-P [C]	2–0–2 [3]
Offered for	B.Tech., M.Tech., Ph.D.	Type	Elective
Prerequisite	Machine Learning or PRML or IML	Antirequisite	None

Objectives

1. To equip students with the ability to design and implement ML/DL systems for real-world applications.
2. To develop expertise in building deployable, reliable, and scalable ML/DL products that meet production standards.

Learning Outcomes

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

1. Design ML/DL systems that solve practical problems and operate in production environments.
2. Diagnose and troubleshoot system faults, identifying root causes effectively.
3. Evaluate and select appropriate frameworks, infrastructures, and trade-offs for scalable deployment.
4. Ensure reproducibility and reliability in ML/DL workflows through experiment tracking and versioning.
5. Analyze trade-offs in system design, balancing accuracy with latency, cost, scalability, and maintainability.
6. Apply techniques for distributed training, mixed precision, and performance optimization.
7. Monitor, update, and maintain deployed ML/DL systems to ensure robustness and fault tolerance.
8. Critically assess emerging trends in ML/DL system design for next-generation applications.

Contents

- Foundations of ML/DL Systems: Data to Deployment Pipeline, Challenges, Lab-to-Production Failures, DevOps to MLOps. (2 lectures)
- Infrastructure & Tools: Virtual Environments, Accelerators, Compilers, Containers, Kubernetes, Reproducibility, Reusability (2 lectures) (2 labs)
- Model Development & Optimization: ML/DL Data Structure and Data Processing, Exploratory Data Analysis and Feature Engineering, Model Development, Hyperparameter Tuning, Model Visualization, Experiment Tracking, Quantization, Mixed Precision, Distributed Training, Large Data Requirement, Inference time optimizations methods like vLLM, TensorRT, ONNX, Model compression methods including pruning and knowledge distillation, libraries and frameworks for model training and inference such as keras, tensorflow, pytorch. (8 lectures) (4 labs)
- Deployment & Operations: Platforms, Versioning, Monitoring, Updating, Troubleshooting, Deployment and Inference Pipelines: CI/CD pipelines, Training Pipeline v/s Inference Pipeline. (4 lectures) (2 labs)
- Performance & Scalability: Trade-offs, Profi ling, Client-Level Inference, Latency Optimization. (4 lectures) (2 labs)
- Advanced Topics: Next-gen DL Systems, Architecture-aware Training, System-level Optimization, LoRA, LLMOps. (4 lectures) (2 labs)
- Ethics & Reliability: Bias, Fairness, Security, Privacy, and Compliance in Production ML Systems, Compliance & Governance: Regulatory requirements (GDPR, HIPAA, AI Act), Explainability. (2 lectures) (1 lab)

The evaluation components should contain a Capstone Project: End-to-End ML/DL System Design with Real-World Deployment with (i) dataset preparation & versioning, (ii) training + experiment tracking, (iii) deployment pipeline, (iv) monitoring + updating. The evaluation components should also include practical assignments in the theory component as well.

Textbooks

1. Online materials and sysML level papers

Self Learning Material

1. https://stanford-cs329s.github.io/
2. https://ucbrise.github.io/cs294-ai-sys-sp19/
3. https://ckaestne.github.io/seai/
4. https://www.coursera.org/specializations/machine-learning-engineering-for-production-mlops#courses
5. MLflow – for MLOps https://mlfl ow.org/docs/latest/ml/
6. Tensor Board – for training visualization https://www.tensorflow.org/tensorboard
7. Jupyterlab notebooks for model training and experimentation. This is also open source https://jupyter.org/.

Title	**Selected Topics in Artificial Intelligence - III**	Number	CSL7XX0
Department	Computer Science and Engineering	L-T-P [C]	3–0–0 [1]
Offered for	B.Tech., M.Tech., PhD	Type	Elective
Prerequisite	Decided by the instructor		

Objectives
The Instructor will:
1. Expose the students to the latest upcoming fields in the area of artificial intelligence

Learning Outcomes
The students are expected to have the ability to:
1. Apply the knowledge of recent topics to specific research areas in the field of artificial intelligence

Contents
The topic clouds for the course include contemporary topics in artificial intelligence and may be updated according to the instructor.

Textbook
Relevant Textbook and/or research papers to be announced by the instructor.

Self-Learning Material
Relevant Textbook and/or research papers to be announced by the instructor.

Preparatory Course Material
Relevant Textbook and/or research papers to be announced by the instructor.

Title	**Computer Graphics (700)**	Number	CSL7450
Department	Computer Science and Engineering	L-T-P [C]	3–0–0 [3]
Offered for	B.Tech, M.Tech., Ph.D.	Type	Elective
Prerequisite	None	Antirequisite	Computer Graphics (400) - CSL4xx

Objectives

The Instructor will:

1. Provide a thorough introduction to computer graphics techniques, focusing on 2D and 3D modelling, image synthesis and rendering

Learning Outcomes

The students are expected to have the ability to:

1. Explain and create interactive graphics application
2. Implement graphics primitives
3. Synthesize and render images for animation and visualization

Contents

CSL7xx1 Introduction to Graphical Primitives 1-0-0 [1]

Introduction: Overview of computer graphics, representing pictures, preparing, presenting & interacting with pictures for presentations; Scan conversion: 2D Geometric Primitives; Area Filling algorithms. Clipping algorithms, Anti Aliasing

Transformations and viewing: 2D and 3D transformations, Matrix representations & homogeneous coordinates, Viewing pipeline, Window to viewport coordinate transformation, clipping operations, viewport clipping, 3D viewing.

CSL7xx2 Graphical Object Representation 1-0-0 [1]

Curves and Surfaces: Conics, parametric and non-parametric forms; Bezier (Bernstein Polynomials) Curves, Cubic-Splines, B-Splines; Quadratic surfaces, Bezier surfaces and NURBS, 3-D modelling.

CSL7xx3 Graphics Rendering 1-0-0 [1]

Hidden surfaces: Depth comparison, Z-buffer algorithm, Back face detection, BSP tree method, the Printer’s algorithm, scan-line algorithm; Hidden line elimination, wire frame methods, fractal - geometry.

Color & shading models: Phong's shading model, Gouraud shading, Shadows and background, Color models, Photo-realistic rendering, Animation and OpenGL primitives: Functions, pipeline, sample programs for drawing 2-D, 3-D objects; event handling and view manipulation, Introduction to GPU and animation

Textbook

1. J. D. Foley, A. Van Dam, S. K. Feiner and J. F. Hughes, Computer Graphics; Principles and practice, Addison Wesley, 2nd Edition in C, 1997.
2. D. F. Rogers and J. A. Adams, Mathematical elements for Computer Graphics, McGraw-Hill, 2nd Edition, 1990.

Self-Learning Material

1. Blender*:* *https://www.blender.org/download/*
2. *OpenGL: http://www.opengl-tutorial.org/*

Preparatory Course Material

1.Department of Computer Science and Engineering, Indian Institute of Technology Madras, https://nptel.ac.in/courses/106106090/

Title	**Systems for Machine Learning**	Number	CSL7XX0
Department	Computer Science and Engineering	L-T-P [C]	3–0–0 [3]
Offered for	PhD, MTech, Executive MTech, UG-4	Type	Elective
Prerequisite	Course in Machine Learning/Deep Learning		

Objectives
The Instructor will:
1. Provide full background of the emerging area lying in the intersection of ML/AI, data management, and systems.
2. Provide knowledge on the landscape and evolution of systems and designs that power modern data science applications, large complex datasets, enterprise analytics, recommendation systems, and social media analytics.

Learning Outcomes

The students are expected to have the ability to:
1. Learn core systems topics that span the entire lifecycle of ML based data analytics.
2. Develop skills for data sourcing and preparation, programming models and systems, scalable ML model building, and fast ML deployment.
3. Build systems that incorporate the correct performance landscape, optimization, and trade-offs between frameworks and compute infrastructure.

Contents

Hardware for Machine Learning
Basics of Computer Organization and Operating Systems, Basics of high performance computing in the cloud, GPU Architecture, GPU Programming Models, GPU Applications design, TPU architecture, Domain Specific Architectures. (10 lectures)
ML Training at Scale
Basics of parallel and scalable data processing, parallelism basics, scalable data access, data parallelism, dataflow systems, classical ML training at scale, deep learning systems, feature engineering and model selection systems, techniques in model compression. (20 lectures)
ML Platforms
Frameworks for Parallel and Distributed Training, DL Compilers, Design of ML Pipelines, Scheduling, ML Deployment, Case studies using Horovod, Distributed PyTorch, TensorFlow Extended (TFX), MLFlow, Michelangelo. (12 lectures)

Textbook
1. Boehm M, Kumar A, Yang J, (2020) Data Management in Machine Learning Systems. Morgan Claypool
2. Ramakrishnan R, Gehrke J, (2002) Data Management Systems (3rd Edition). Irwin Computer Science.

References/Self-Learning Material
1. Computer Organization and Design (5th edition), by David Patterson and John Hennessy
2. Programming Massively Parallel Processors, (3rd Edition) by Wen Mei Hwu and David B Kirk.
3. ML Framework guides from respective websites.
4. Research papers.