

## Computer Science Engineering

SL. No	Old Course code		Name of Course	L-T-P-C	Proposed Level (UG/PG)
1	CS 601	CS602T	<a href="#">Software Development for Scientific Computing</a>	3-0-0-6	PG
2	CS 603	CS603T	<a href="#">Approximation algorithms</a>	3-0-0-6	PG
3	CS 604	CS604T	<a href="#">Parametrized Algorithms and Complexity</a>	3-0-0-6	PG
4	CS 605	CS605T	<a href="#">Reinforcement Learning</a>	3-0-0-6	PG
5	CS 606	CS502T	<a href="#">Advanced Topics in Embedded Computing</a>	3-0-0-6	PG
6	CS 607	CS606T	<a href="#">Advanced Computer Networks</a>	3-0-0-6	PG
7	CS 608	CS607T	<a href="#">FPGA for communication networks prototyping</a>	3-0-0-6	PG
8	CS 609	CS608T	<a href="#">Software Defined Networking (SDN) and Network Function Virtualization (NFV)</a>	3-0-0-6	PG
9	CS 610	CS609T	<a href="#">Advanced Distributed Systems</a>	3-0-0-6	PG
10	CS 611	CS605C	<a href="#">Advanced Software Systems Lab</a>	0-1-6-8	PG
11	CS 612	CS501L	<a href="#">Statistical Pattern Recognition Laboratory</a>	0-0-3-3	PG
12	CS 614	CS601L	<a href="#">Reinforcement Learning Laboratory</a>	0-0-3-3	PG
13	CS 616	CS503T	<a href="#">Statistical Pattern Recognition</a>	3-0-0-6	PG
14	CS 617	CS610T	<a href="#">Special Topics in Hardware Systems</a>	3-0-0-6	PG
15	CS 620	CS611T	<a href="#">Formal Models for Concurrent and Asynchronous systems</a>	3-0-0-6	PG
16	CS 621	CS612T	<a href="#">Logic and Applications</a>	3-0-0-6	PG
17	CS 622	CS613T	<a href="#">Special Topics in Automata and Logics</a>	3-0-0-6	PG
18	CS 623	CS614T	<a href="#">Advanced Topics in Communication Networks</a>	3-0-0-6	PG
19	CS 624	CS620T	<a href="#">Compilers - Principles and Implementation</a>	3-0-0-6	PG
20	CS 625		<a href="#">Topics in Stochastic Control and Reinforcement Learning</a>	3-0-2-8	PG
21	CS 626		<a href="#">Topics in Data Structures and Algorithms</a>	2-0-2-6	PG
22	CS 627		<a href="#">Data Structures</a>	3-0-0-6	PG
23	CS 628		<a href="#">Algorithms</a>	3-0-0-3	PG
24	CS 629		<a href="#">Introduction to Reinforcement Learning</a>	2-0-2-6	PG
25	CS 630		<a href="#">Statistical Machine Learning</a>	2-0-2-6	PG
26	CS 631	CS601S	<a href="#">Seminar</a>	0-0-4-4	PG

27	CS 632	CS504T	<a href="#">Runtime Verification</a>	3-0-0-6	PG
28	CS 702	CS601C	<a href="#">Systems Bootcamp for ML</a>	1-0-2-4	PG
29	CS 703		<a href="#">Topics in Data Structures and Algorithms</a>	3-0-0-6	PG
30	CS 704		<a href="#">Advanced Algorithms</a>	3-0-0-6	PG
31	CS 705		<a href="#">Topics in Graph Theory</a>	3-0-0-6	PG
32	CS 706		<a href="#">Topics in Parameterized Algorithms and Complexity</a>	3-0-0-6	PG
33	CS 801	CS505T	<a href="#">Power Aware Computing</a>	3-0-2-8	PG
34	CS 802		<a href="#">Dataflow Processor Architecture (Guided Study)</a>	3-0-0-6	PG
35	CS 810	CS501C	<a href="#">Advanced Computer Architecture</a>	3-0-3-9	PG
36	CS 438	CS615T	<a href="#">Natural Language Processing</a>	3-0-0-6	PG
37	CS 439	CS412T	<a href="#">Introduction to Sanskrit Computational Linguistics</a>	3-0-0-6	PG
38	CS 440	CS413T	<a href="#">Advances in Cloud Technologies</a>	1-0-0-2	PG
39	CS 636	CS616T	<a href="#">Advanced Data Structures and Algorithms Lab</a>	0-0-3-3	PG
40	CS 633	CS618T	<a href="#">Advanced Data Structures and Algorithms Lab</a>	3-0-0-6	PG
41	CS 618	CS602C	<a href="#">Advanced Software Development Laboratory</a>	1-0-4-6	PG
42	CS 638	CS617T	<a href="#">Combinatorics and Probability</a>	3-0-0-6	PG
43	CS 639	CS619T	<a href="#">Scalable Data Mining</a>	3-0-0-6	PG

44	CS 101	CS101C	<a href="#">Computer Programming</a>	3-0-2-8	UG
45	CS 103		<a href="#">Introduction to Programming – 1</a>	3-0-2-4	UG
46	CS 104		<a href="#">Introduction to High Performance Computing</a>	0.5-0-0-1	UG
47	CS 106		<a href="#">Data Structures and Algorithms</a>	3-0-0-6	UG
48	CS 111	CS101L	<a href="#">Data Structures and Algorithms Laboratory</a>	0-0-3-3	UG
49	CS 202	CS201T	<a href="#">Automata Theory</a>	3-1-0-8	UG
50	CS 203	CS202T	<a href="#">Discrete Structures</a>	3-0-0-6	UG
51	CS 205	CS203T	<a href="#">Design and Analysis of Algorithms</a>	3-0-0-6	UG
52	CS 209	CS204T	<a href="#">Artificial Intelligence</a>	3-0-0-6	UG
53	CS 213	CS201C	<a href="#">Software Systems Laboratory</a>	1-3-0-8	UG
54	CS 214	CS201L	<a href="#">Artificial Intelligence Lab</a>	0-0-3-3	UG

55	CS 301	CS302T	<a href="#">Computer Architecture</a>	3-0-0-6	UG
56	CS 303	CS202L	<a href="#">Data Bases and Information Systems</a>	3-0-0-6	UG
57	CS 304	CS303T	<a href="#">Operating Systems</a>	3-0-0-6	UG
58	CS 306	CS301T	<a href="#">Introduction to Artificial Neural Networks</a>	3-0-0-6	UG
59	CS 309	CS301P	<a href="#">Research and Development Project</a>	6 credits	UG
60	CS 311	CS301L	<a href="#">Computer Architecture Laboratory</a>	0-0-3-3	UG
61	CS 313	CS202L	<a href="#">Data Bases and Information Systems Laboratory</a>	0-0-3-3	UG
62	CS 314	CS302L	<a href="#">Operating Systems Laboratory</a>	0-0-3-3	UG
63	CS 315	CS303L	<a href="#">Computer Networks Laboratory</a>	0-0-3-3	UG
64	CS 316	CS304L	<a href="#">Compilers Lab</a>	0-0-3-3	UG
65	CS 320	CS302P	<a href="#">Research and Development Project II</a>	6 credits	UG
66	CS 323	CS304T	<a href="#">Compilers</a>	3-0-0-6	UG
67	CS 324		<a href="#">Programming Techniques</a>	3 credits	UG
68	CS 348	CS305T	<a href="#">Computer Networks</a>	3-0-0-6	UG
69	CS 401	CS401T	<a href="#">Software Engineering</a>	0-0-3-3	UG
70	CS 402	CS402T	<a href="#">Distributed Systems</a>	3-0-0-6	UG
71	CS 403	CS403T	<a href="#">Graph Theory and Combinatorics</a>	3-0-0-6	UG
72	CS 405	CS401P	<a href="#">B.Tech Project-CSE</a>	6 credits	UG
73	CS 409	CS402P	<a href="#">B.Tech Project-CSE - II</a>	6 credits	UG
74	CS 410	CS404T	<a href="#">Parallel Computing</a>	3-0-0-6	UG
75	CS 421		<a href="#">Logic for Computer Science</a>	3-0-0-6	UG
76	CS 422	CS406T	<a href="#">Principles of Programming Languages</a>	3-0-0-6	UG
77	CS 424		<a href="#">Elements of Programming</a>	2-0-2-6	UG
78	CS 426	CS501T	<a href="#">Introduction to Blockchains</a>	3-0-0-6	UG
79	CS 427	CS407T	<a href="#">Mathematics for Data Science</a>	3-0-0-6	UG
80	CS 428	CS601T	<a href="#">Deep Learning</a>	3-0-0-6	UG
81	CS 430	CS401C	<a href="#">Computer Graphics</a>	3-0-2-8	UG
82	CS 432	CS408T	<a href="#">Introduction to Logic</a>	3-0-0-6	UG
83	CS 433	CS409T	<a href="#">Cloud Software Development</a>	3-0-0-3	UG
84	CS 434	CS410T	<a href="#">Programming Parallel Machines</a>	3-0-0-6	UG
85	CS 435	CS411T	<a href="#">Introduction to Abstractions and Paradigms for Programming</a>	1-0-0-2	UG
86	CS 436		<a href="#">Operational Analysis</a>	3-0-0-6	UG

87	CS 437	CS506T	<a href="#">Guided Study</a>	6 credits	UG
88	CS440		<a href="#">Advanced in Cloud Technologies</a>	1-0-0-2	UG
89	CS 707	CS603C	<a href="#">Advanced Topics in Deep learning</a>	3-0-2-8	PG
90	CS 708	CS701T	<a href="#">Adversarial Machine Learning</a>	3-0-0-6	PG
91	CS 709	CS701C	<a href="#">Machine Learning Applications in Wireless Networks</a>	2-1-0-6	PG
92	CS621T	CS621T	<a href="#">Cryptography and Network Security</a>	2-1-0-6	PG
93		CS603L	<a href="#">Scalable Data Mining Lab</a>	0-0-3-3	PG
94		CS604C	<a href="#">Introduction to Open Radio Access Networks</a>	2-1-0-6	PG
95		CS602L	<a href="#">Advanced Data Structures and Algorithms Lab</a>	0-0-3-3	PG
96		CS301P	<a href="#">Research and Development Project</a>	0-0-6-6	
97		CS701P	<a href="#">MTech Technical Project – I</a>	32	
98		CS702P	<a href="#">MTech Technical Project – II</a>	32	
99		CS101T	<a href="#">Data Structures and Algorithms</a>	3-0-0-6	UG
100		CS205T	<a href="#">Data Bases and Information Systems</a>		
101		CS405T	<a href="#">Logic for Computer Science</a>	3-0-0-6	UG

1	<b>Title of the course (L-T-P-C)</b>	<b>Software Development for Scientific Computing (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Data Structures and Algorithms, C / C++ / Java / Matlab
3	<b>Course content</b>	Algorithmic Patterns in Scientific Computing: dense and sparse linear algebra, structured and unstructured grid methods, particle methods (N-body, Particle-Particle, Particle-in-cell, Particle-in-a- mesh), Fast Fourier Transforms, Implementing PDEs, C++ standard template library (STL), Introduction to debugging using GDB, GMake, Doxygen, Version Control System, Profiling and Optimization, asymptotic analysis and algorithmic complexity. Mixed-language programming using C, Fortran, Matlab, and Python, Performance analysis and high-performance code, Data locality and auto tuning, Introduction to the parallel programming world.
4	<b>Texts/References</b>	<ul style="list-style-type: none"> <li>● Stroustrup C++ Language Reference (<a href="https://www.stroustrup.com/4th.html">https://www.stroustrup.com/4th.html</a>)</li> <li>● Suely Oliveira, David Steward: Writing Scientific Software: A Guide to Good Style. Cambridge University Press, 2006</li> <li>● Web references to GNU Make, GDB, Git, GProf, Gcov.</li> <li>● Code Complete: A Practical Handbook of Software Construction</li> <li>● <a href="https://www2.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-183.html">https://www2.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-183.html</a></li> </ul>

1	<b>Title of the course</b> (L-T-P-C)	<b>Approximation algorithms</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Data Structures and Algorithms (CS201)
3	<b>Course content</b>	Introduction, approximation schemes, design and analysis of approximation algorithms - combinatorial algorithms, linear programming based algorithms. Hardness of approximation.
4	<b>Texts/References</b>	<b>Textbook:</b> Approximation algorithms. Vazirani, Vijay V. Berlin: springer, 2001. <b>Reference:</b> The design of approximation algorithms. Williamson, David P., and David B. Shmoys. Cambridge university press, 2011.

1	<b>Title of the course (L-T-P-C)</b>	<b>Parametrized Algorithms and Complexity (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Data Structures and Algorithms, Design and Analysis of Algorithms
3	<b>Course content</b>	Introduction. Kernelization, Bounded Search Trees, Iterative Compression, Treewidth, Advanced kernelization algorithms. Lower bounds: Fixed- parameter intractability, lower bounds based on ETH, lower bounds for kernelization. <b>Parameterized Algorithms, Kernelization, and Complexity of Graph Modification Problems</b>
4	<b>Texts/References</b>	<b>Textbook:</b> <ul style="list-style-type: none"> <li>• Parameterized Algorithms, Marek Cygan, Fedo</li> <li>• V. Fomin, Lukasz Kowalik. Daniel Lokshtanov, Daniel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Sourabh. Springer. 2015</li> </ul> <b>Reference:</b> <ul style="list-style-type: none"> <li>• Parameterized Complexity, R. G. Downey, and M. R. Fellows. Springer Science and Business Media. 2012</li> </ul>

1	<b>Title of the course</b> (L-T-P-C)	<b>Reinforcement Learning</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Basic Probability and Linear Algebra
3	<b>Course content</b>	Bandit Algorithms -- Regret based - UCB, Thomson Sampling, PAC Based - Median Elimination, Markov Decision Process Modeling - Bellman Equation, Dynamic Programming Solutions - Value and Policy Iteration, Linear Programming, Model free methods - Monte Carlo and Temporal Difference Methods - Q-learning, Value function Approximation - State Aggregation, Critic Only/Value Based Methods Methods - TD methods, Q- Learning, SARSA, Actor Only/Policy Based methods - Reinforce, Actor-Critic Methods - Policy Gradient, Natural Actor Critic, Deep RL - DQN, A3C, Model Based RL, Integrating Learning and Planning, Case-studies.
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Richard S. Sutton and Andrew G. Barto, Introduction to Reinforcement Learning, 2nd Edition, MIT Press. 2017.</li> <li>2. Dimitri Bertsekas and John G. Tsitsiklis, Neuro Dynamic Programming, Athena Scientific. 1996.</li> <li>3. Bertsekas, Dimitri P. Dynamic Programming and Optimal Control. Vol. 1 and 2. 4th edition, 2012.</li> <li>4. Algorithms for Reinforcement Learning, Csaba Szepesvári, Morgan &amp; Claypool, 2009.</li> <li>5. Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems, Sébastien Bubeck and Nicolò Cesa-Bianchi, Foundations and Trends in Machine Learning, Volume 5, Number 1, 2012.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Reinforcement Learning 2</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	CS 301
3	<b>Course content</b>	Introduction to systems software in embedded platforms Boot loader Embedded Linux kernel (Processes, Threads, Interrupts) Device Drivers Scheduling Policies (including Real Time) Memory Management Optimizations (Data level and Memory level) Embedded Systems Security Introduction to Embedded GPUs and Accelerators Embedded Heterogenous Programming with Open CLApplication Case Study on Embedded Platforms – eg. Neural Network inferencing on Embedded Platforms, Advanced Driver Assistance Systems
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Building Embedded Linux Systems, 2nd Edition by Gilad Ben-Yossef, Jon Masters, Karim Yaghmour, Philippe Gerum, O'Reilly Media, Inc. 2008</li> <li>2. Linux Device Drivers, Third Edition By Jonathan Corbet, Alessandro Rubini, Greg Kroah-Hartman, O'Reilly Media, Inc. 2005</li> <li>3. Embedded Systems: ARM Programming and Optimization by Jason D Bakos, Elsevier, 2015</li> <li>4. Learning Computer Architecture with Raspberry Pi by Eben Upton, Jeff Duntemann, Ralph Roberts, Tim Mamtora, Ben Everard, Wiley Publications, 2016</li> <li>5. Real Time Systems by Jane S. Liu, 1 edition, Prentice Hall; 2000</li> <li>6. Practical Embedded Security: Building Secure Resource-Constrained Systems by Timothy Stapko, Elsevier, 2011</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Advanced Computer Networks</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	
3	<b>Course content</b>	<ol style="list-style-type: none"> <li>1. Circuit, Packet and Virtual Circuit Switching, MPLS</li> <li>2. Switch Architectures, Buffering Strategies, Input and Output Queuing, IP Buffer Sizing</li> <li>3. Quality of Service and Scheduling Algorithms</li> <li>4. IP Address Lookup and IP Packet Classification algorithms</li> <li>5. Software Defined Networking</li> <li>6. Next Generation Network Architectures, Network Provisioning and Design, and “Green” (Energy- Efficient) Networking</li> <li>7. Data Driven Networking</li> <li>8. Wireless Networks - MANETs, Sensor Networks, Cellular Networks, Personal Area Networks</li> <li>9. Content Based Delivery Networks - Principles of data dissemination, aggregation and caching that are applied to sensor networks, Internet of Things, and other content-based paradigms. Students will survey recent research publications on opportunistic networks and next generation content-based networking ideas.</li> <li>10. Delay tolerant Networks</li> <li>11. Network security - authentication, access control, privacy preservation, intrusion detection and prevention</li> <li>12. Performance analysis of new Networking ideas using simulation (such as Network Simulator (ns3), GENI testbed, Simulink, Open LTE and Open C-RAN frameworks)</li> </ol>
4	<b>Texts/References</b>	<p><b>Textbook:</b>  Computer Networks: A Systems Approach, Larry Peterson and Bruce Davie, 2011.  Performance Evaluation of Computer Systems, by Raj Jain, Wiley, 1991.  Computer Networking, Kurose and Ross, Addison-Wesley, 2012.</p> <p><b>Reference:</b></p> <ol style="list-style-type: none"> <li>1. An Engineering Approach to Computer Networking by S. Keshav, 1997, Addison-Wesley Professional Series.</li> <li>2. Network Routing, by Deepankar Medhi and Karthikeyan Ramasamy, Morgan Kaufmann, 2007.</li> <li>3. SDN: Software Defined Networks, by Thomas D. Nadeau, Ken Gray, O’Reilly Media, 2013.</li> <li>4. High Performance Switches and Routers, By H.Jonathan Chao and Bin Liu, Wiley, 2007. Network Algorithmics, by George Varghese, Morgan Kaufmann, 2005</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>FPGA for communication networks prototyping (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	EE 224 Digital System Exposure on Computer Network
3	<b>Course content</b>	History and evaluation of FPGAs; FPGA architecture; Introduction to Quartus Prime (vendors and design tools; vendors and programmable logic); Exploiting Simulation tools (e.g., ModelSim); Exploiting FPGAs for multi-domain technologies; Introduction to radio access networks-fronthaul (e.g., common public radio interface); optical network; metro and core networks; Cross-layer design; The role of FPGA in the specified network segments and use case scenarios; In and Out; Clocks and Registers; State Machines; Modular Design; Memories Managing Clocks; I/O Flavors; Exploiting Qsys and Nios II tools.
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. C. Maxfield, "The Design Warrior's Guide to FPGAs: Devices, Tools and Flows", Jun. 2004, eISBN 9780080477138</li> <li>2. FPGAs For Dummies, 2nd Intel Special Edition. Published by. John Wiley &amp; Sons, Inc</li> <li>3. William J. Dally, R. Curtis Harting, "Digital Design: A Systems Approach 1st Edition", Cambridge University Press, September 2012, ISBN 9780521199506</li> <li>4. Verilog by Example: A Concise Introduction for FPGA Design, Blaine C. Readler</li> <li>5. Course materials: Slides; Notes; Tutorials from Altera website <a href="https://www.altera.com/support/training/university/materials-tutorials.html">https://www.altera.com/support/training/university/materials-tutorials.html</a></li> <li>6. R. Ramaswami, K. Sivarajan, G. Sasaki; "Optical Networks: A Practical Perspective," 3rd Ed., Morgan Kaufmann, ISBN: 9780123740922.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Software Defined Networking (SDN) and Network Function Virtualization (NFV)</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Computer Networks
3	<b>Course content</b>	<p>History and evolution of SDN; SDN Architecture (Application, Control, Infrastructure Layer); SDN Interfaces (East/West/North/South-bound interfaces); SDN Security; SDN routing; SDN standards; SDN Controllers; Network Operating Systems and Languages; OpenFlow; Software Switches (e.g. OpenVSwitch); SDN Simulation/Emulation Platforms (e.g. Mininet); Federated SDN networks; SDN Applications and Use Cases; Programming assignment/project;</p> <p>Need for NFV; NFV and SDN Relationship; Virtual Network Functions; Service Function Chaining; NFV Specifications; NFV Architecture; NFV Use Cases; NFV Management and orchestration (MANO); Open-source NFV; Hands-on exercises based on OpenStack/Docker.</p>
4	<b>Texts/References</b>	<ul style="list-style-type: none"> <li>● Software Defined Networks: A Comprehensive Approach by Paul Goransson and Chuck Black, Morgan Kaufmann Publications, 2014</li> <li>● SDN – Software Defined Networks by Thomas D. Nadeau &amp; Ken Gray, O'Reilly, 2013</li> <li>● Software Defined Networking with OpenFlow, By Siamak Azodolmolky, Packt Publishing, 2013</li> <li>● Gray, Ken, and Thomas D. Nadeau. Network function virtualization. Morgan Kaufmann, 2016.</li> <li>● Zhang Ying. Network Function Virtualization: Concepts and Applicability in 5G Networks. John Wiley &amp; Sons, 2018.</li> <li>● Foundations of modern networking- SDN, NFV, QoE, IoT, and Cloud, William Stallings</li> <li>● James Kurose and Keith Ross, "Computer Networking, A Top-Down Approach"</li> </ul>

1	<b>Title of the course</b> (L-T-P-C)	<b>Advanced Distributed Systems</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Operating Systems, Data Structures and Algorithms, Programming in C++
3	<b>Course content</b>	<p>Synchronization, Global Snapshot and Distributed Mutual Exclusion, Consensus &amp; Agreement, Checkpointing &amp; Rollback Recovery, Deadlock Detection, Termination Detection, Message Ordering &amp; Group Communication, Fault Tolerance and Self-Stabilization, Peer to Peer Systems</p> <p>Mining Data Streams in a distributed system: filtering data streams, queries on streams, pattern detection</p> <p>Key-Value Storage: Cassandra, HBase</p> <p>Virtualization and Cloud Computing: virtual machines containers Message oriented communication, Publish Subscribe Systems (use case Apache Kafka)</p> <p>Security: Distribution of security mechanisms, access control, and security management.</p>
4	<b>Texts/References</b>	<p>Distributed Computing: Principles, Algorithms, and Systems- Ajay D. Kshemkalyani and Mukesh Singhal</p> <p>Mining Massive data sets- Jure Leskovec, Anand Rajaraman, Jeff Ullman</p> <p>Distributed Algorithms – An Intuitive Approach (The MIT Press) by Wan Fokkink</p> <p>Distributed Algorithms-Nancy Lynch</p>

1	<b>Title of the course</b> (L-T-P-C)	<b>Advanced Software Systems Lab</b> <b>(0-1-6-8)</b>
2	<b>Pre-requisite courses(s)</b>	Nil
3	<b>Course content</b>	The course should involve 3 large programming projects. A set of three project oriented assignments which will be announced at the start of each semester with definite submission deadlines. The set of assignments will be designed to develop skills and familiarity with a majority of the following: make, configuration management tools, installation of software, archiving and creation of libraries, version control systems, documentation and literate programming systems, GUI creation, distributed state maintenance over a network, programming in different environments like desktop and handhelds, program parsing and compilation including usage of standard libraries like pthreads, numerical packages, XML and semi-structured data, simulation environments, testing and validation tools.
4	<b>Texts/References</b>	<ul style="list-style-type: none"> <li>• Unix concepts and applications by Sumitabha Das</li> <li>• PThreads Programming A POSIX Standard for Better Multiprocessing By Dick Buttlar, Jacqueline Farrell, Bradford Nichols</li> <li>• Head-First Python by Paul Barry</li> </ul>

1	<b>Title of the course</b> (L-T-P-C)	<b>Statistical Pattern Recognition Laboratory</b> <b>(0-0-3-3)</b>
2	<b>Pre-requisite courses(s)</b>	Currently taking statistical pattern recognition theory course
3	<b>Course content</b>	The lab will closely follow the theory course. The idea is to have the students implement the basic algorithms on different topics studied in the statistical pattern recognition theory course.
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. R.O.Duda, P.E.Hart and D.G.Stork, Pattern Classification, John Wiley, 2001.</li> <li>2. C.M.Bishop, Pattern Recognition and Machine Learning, Springer, 2006.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Reinforcement Learning Laboratory</b> <b>(0-0-3-3)</b>
2	<b>Pre-requisite courses(s)</b>	Currently taking reinforcement learning theory course
3	<b>Course content</b>	The lab will closely follow the theory course. The idea is to have the students implement the basic algorithms on different topics studied in the reinforcement learning theory course.
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Richard S. Sutton and Andrew G. Barto, Introduction to Reinforcement Learning, 2nd Edition, MIT Press. 2017.</li> <li>2. Dimitri Bertsekas and John G. Tsitsiklis, Neuro Dynamic Programming, Athena Scientific. 1996.</li> <li>3. Bertsekas, Dimitri P. Dynamic Programming and Optimal Control. Vol. 1 and 2. 4th edition, 2012.</li> <li>4. Algorithms for Reinforcement Learning, Csaba Szepesvári, Morgan &amp; Claypool, 2009.</li> <li>5. Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems, Sébastien Bubeck and Nicolò Cesa-Bianchi, Foundations and Trends in Machine Learning, Volume 5, Number 1, 2012.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Statistical Pattern Recognition</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Multivariate Calculus and Linear Algebra, Probability, Programming
3	<b>Course content</b>	Bayesian Decision Making and Bayes Classifier, Parametric and Non Parametric Estimation of Densities, General Linear Models, Discriminative Learning based Models, Dimensionality Reduction Techniques, Empirical and Structural risk minimization, Ensemble Methods - Bagging, Boosting, Pattern Clustering, Graphical Models, Statistical Learning Theory.
4	<b>Texts/References</b>	<ul style="list-style-type: none"> <li>• R.O.Duda, P.E.Hart and D.G.Stork, Pattern Classification, John Wiley, 2001.</li> <li>• C.M.Bishop, Pattern Recognition and Machine Learning, Springer, 2006.</li> </ul>

1	<b>Title of the course</b> (L-T-P-C)	<b>Special Topics in Hardware Systems</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Digital Systems Computer Architecture
3	<b>Course content</b>	Introduction Processors Memories Special Processors and Accelerations Architecture for Machine Learning Detailed Syllabus is attached in the appendix
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. J .L. Hennessy, D. A. Patterson: Computer Architecture :A Quantitative Approach(The Morgan Kaufmann Series in Computer Architecture and Design),2011</li> <li>2. J. P. Shen, M. H. Lipasti: Modern Processor Design: Fundamentals of Superscalar Processors , Waveland Press, 2013</li> <li>3. B. Reagan, R. Adolf, P. Whatmough, G.Y-Wei, D.Brooks:Deep Learning for Computer Architects Synthesis Lectures on Computer Architecture, Morgan &amp; Claypool,2017.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Logic and Applications</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Discrete Mathematics, Theory of computation
3	<b>Course content</b>	<p><b>Module 1:</b>Propositional Logic:Natural deduction, semantics, soundness, completeness, compactness, normal forms, Horn clauses and satisfiability.</p> <p><b>Module 2:</b> Predicate Logic: Natural deduction, resolution, undecidability, expressiveness.</p> <p><b>Module 3:</b> Some decidable fragments of first-order logic and their decision procedures: propositional logic, equality with uninterpreted functions, linear arithmetic, Presburger logic ,bit vectors, arrays, pointer logic.</p> <p><b>Module 4:</b> SAT and SMT solvers: theory and practice: Decision procedures for combinations of first-order theories: Nelson-Oppen, Shostak, Satisfiability Modulo Theories(SMT)Combination with SAT solvers: eager, lazy approaches. Student is required to do a small project using a SAT/SMT solver.</p>
4	<b>Texts/References</b>	<ul style="list-style-type: none"> <li>• Logic in Computer Science, Michael Huth and Mark Ryan, Cambridge University Press.</li> <li>• Mathematical Logic for Computer science, Mordechai Ben-Ari, Springer.</li> <li>• Logic for Computer Scientists, Uwe Schoning, Birkhauser.</li> <li>• SAT/SMT by example, Dennis Yurichev.</li> </ul>

1	<b>Title of the course</b> (L-T-P-C)	<b>Special Topics in Automata and Logics</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Discrete Mathematics, Theory of computation, Logic and its applications.
3	<b>Course content</b>	<p>This course aims at giving an introduction to the theory of automata working on infinite words and infinite trees and connections thereof to logics. These automata and related logics are of fundamental importance in the areas of formal specification and verification of reactive systems. If time permits, we will also discuss some basic results in finite model theory. Below is a list of topics, which will be discussed in this course. Automata on finite words - equivalence of MSO and automata; Automata on infinite words – different acceptance conditions; Closure properties and equivalence of different acceptance conditions and related translations Determinization and complementation results; Equivalence of automata and MSO and decidability of MSO; Automata on infinite trees - different acceptance conditions; Closure properties and comparison of expressive power of different acceptance conditions and related translations Complementation result for tree automata via parity games; Equivalence of MSO and tree automata; Decidability of MSO over trees; Parity games and determinacy; Ehrenfeucht-Fraïssé games in logics and applications.</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Wolfgang Thomas: Automata on infinite objects, Handbook of theoretical computer science (vol B): formal methods and semantics, Elsevier.</li> <li>2. Wolfgang Thomas: Languages, automata, and logic, Handbook of formal languages, vol. 3: beyond words, Springer-Verlag.</li> <li>3. Dominique Perrin, Jean-Eric Pin: Infinite words, Elsevier</li> <li>4. Erich Gradel, Wolfgang Thomas, Thomas Wilke: Automata, logics, and infinite games: a guide to current research. LNCS, Springer-Verlag.</li> <li>5. Leonid Libkin: Elements of finite model theory, Springer-Verlag.</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Advanced Topics in Communication Networks (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Undergraduate Computer Networks course, Good Programming Background.
3	<b>Course content</b>	<ol style="list-style-type: none"> <li>1. 4G/5G Networks – Radio Access Architecture, Evolved Packet Core, Protocols, Network Management Algorithms, Network Optimization, Resource Allocation Algorithms, Security.</li> <li>2. Fog Computing, Edge Computing – Architecture, Optimization, Resource Allocation, and Load Balancing</li> <li>3. Internet of Things</li> <li>4. Data Driven Networking</li> <li>5. Application of SDN and NFV in next generation IoT/cellular networks.</li> </ol>
4	<b>Texts/References</b>	Research papers and online courses from Coursera/Udacity/Edx will be referred to for learning the afore-mentioned topics.

1	<b>Title of the course (L-T-P-C)</b>	<b>Compilers - Principles and Implementation (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Data Structures and Algorithms, Computer Architecture, Automata Theory
3	<b>Course content</b>	Structure of a compiler, the compiled and interpreted execution models. Lexical analysis and parsing using lex and yacc. Scope and visibility analysis. Data layout and lifetime management of data. Runtime environment. Dynamic memory allocation and Garbage collection. Translation of expressions, control structures, and functions. Code generation and local optimizations, Lattice theory, register allocation, instruction scheduling, optimizations - dataflow, control flow, reaching definitions, and liveness analysis, code transformation-tiling, fusion.
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Alfred V. Aho, Monica S. Lam, Ravi Sethi and Jeffrey D.Ullman: Compilers: Principles, Techniques, and Tools, 2/E, AddisonWesley 2007.</li> <li>2. Andrew Appel: Modern Compiler Implementation in C/ML/Java, Cambridge University Press, 2004</li> <li>3. Dick Grune, Henri E. Bal, Cerial J.H. Jacobs and Koen G. Langendoen: Modern Compiler Design, John Wiley &amp; Sons, Inc. 2000.</li> <li>4. Michael L. Scott: Programming Language Pragmatics, Morgan Kaufman Publishers, 2006.</li> <li>5. Fisher and LeBlanc: Crafting a Compiler in C.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Topics in Stochastic Control and Reinforcement Learning</b> <b>(3-0-2-8)</b>
2	<b>Pre-requisite courses(s)</b>	Probability, Linear Algebra and Multi-variable Calculus
3	<b>Course content</b>	Bandit Algorithms -- Regret based - UCB, Thomson Sampling, PAC Based - Median Elimination, Optimality of Bandit Algorithms, Finite Horizon Problems, Infinite Horizon Problems - Total Cost Criterion, Discounted Cost Criterion, Average Cost Criterion, Markov Decision Process (MDP), Partially observable MDP (POMDP), and Dynamic Programming Solutions - Value and Policy Iteration, Model free methods - Monte Carlo and Temporal Difference Methods - Q-learning, SARSA, on/off-policy learning, Stochastic Approximation: Single and multi-timescale stochastic approximation, ordinary differential equation based convergence results. Convergence of SARSA and Q- learning, Value function Approximation - State Aggregation, Critic Only/Value Based Methods - TD methods, gradient temporal difference learning, Actor Only/Policy Based methods - Reinforce, Actor-Critic Methods - Policy Gradient, Natural Actor Critic, Deep RL - DQN, A3C, Regret Based RL - Upper Confidence Reinforcement Learning, Posterior sampling Reinforcement Learning.
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Bertsekas, Dimitri P. Dynamic Programming and Optimal Control. Vol. 1 and 4th edition, 2012.</li> <li>2. Algorithms for Reinforcement Learning, Csaba Szepesvári, Morgan &amp; Claypool, 2009.</li> <li>3. Dimitri Bertsekas and John G. Tsitsiklis, Neuro Dynamic Programming, Athena Scientific. 1996.</li> <li>4. Richard S. Sutton and Andrew G. Barto, Introduction to Reinforcement Learning, 2nd Edition, MIT Press. 2017.</li> <li>5. Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems, Sébastien Bubeck and Nicolò Cesa-Bianchi, Foundations and Trends in Machine Learning, Volume 5, Number 1, 2012.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Topics in Data Structures and Algorithms</b> <b>(2-0-2-6)</b>
2	<b>Pre-requisite courses(s)</b>	No prerequisites
3	<b>Course content</b>	Data structures - arrays, linked lists, stacks and queues, heap and binary search tree. Algorithm design techniques - divide and conquer, greedy, and dynamic programming. Algorithms for graph problems. Asymptotic notations, Complexity lower bounds and NP-completeness.
4	<b>Texts/References</b>	<p><b>Textbook:</b></p> <ol style="list-style-type: none"> <li>1. Cormen, Leiserson, Rivest and Stein, <b>Introduction to Algorithms</b>, 3rd edition, by, MIT Press, 2009.</li> </ol> <p><b>Reference:</b></p> <ol style="list-style-type: none"> <li>1. Sanjoy Dasgupta, Christos Papadimitriou and Umesh Vazirani, <b>Algorithms</b>, McGraw Hill Education, 2008.</li> <li>2. Kleniberg and Tardos, <b>Algorithm Design</b>, 1st edition, Pearson, 2006.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Data Structures</b> <b>(3-0-0-3)</b>
2	<b>Pre-requisite courses(s)</b>	None
3	<b>Course content</b>	<ul style="list-style-type: none"> <li>• Asymptotic notations-Theta, Big-oh, Big-omega, little-oh, little-omega</li> <li>• Data structures - arrays, linked lists, stacks and queues, heap and binary search tree, hash tables</li> </ul>
4	<b>Texts/References</b>	<p><b>Textbook:</b></p> <ol style="list-style-type: none"> <li>1. Cormen TH, Leiserson CE, Rivest RL, Stein C. Introduction to algorithms. MIT press; 2009.</li> </ol> <p><b>Reference:</b></p> <ol style="list-style-type: none"> <li>1. Brass P. Advanced data structures. Cambridge: Cambridge University Press; 2008.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Algorithms</b> <b>(3-0-0-3)</b>
2	<b>Pre-requisite courses(s)</b>	
3	<b>Course content</b>	Algorithm design techniques - divide and conquer, greedy, and dynamic programming, Algorithms for graph problems. Complexity, lower bounds, and NP-completeness.
4	<b>Texts/References</b>	<b>Textbook</b> <ol style="list-style-type: none"> <li>1. Cormen TH, Leiserson CE, Rivest RL, Stein C. Introduction to algorithms. MIT press; 2009.</li> </ol> <b>Reference:</b> <ol style="list-style-type: none"> <li>1. Dasgupta S, Papadimitriou CH, Vazirani UV. Algorithms. New York: McGraw-Hill Higher Education; 2008..</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Introduction to Reinforcement Learning (2-0-2-6)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Probability Models, Stochastic Process and its Applications, Math for Data Science I and II
3	<b>Course content</b>	<p>andit Algorithms -- Regret based - UCB, Thomson sampling, Markov Decision Process Modeling - Bellman equation, Dynamic Programming Solutions - Value and policy Iteration, Model free methods - Monte Carlo and temporal Difference Methods - Q-learning, Value function Approximation - State Aggregation, Critically/Value Based Methods Methods - TD methods, -Learning, SARSA, Actor Only/Policy Based methods - reinforce, Actor-Critic Methods - Policy Gradient, Deep L - DQN, A3C.</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Richard S. Sutton and Andrew G. Barto, Introduction to Reinforcement Learning, 2nd Edition, MIT Press. 2017.</li> <li>2. Dimitri Bertsekas and John G. Tsitsiklis, Neuro Dynamic Programming, Athena Scientific. 1996.</li> <li>3. Bertsekas, Dimitri P. Dynamic Programming and Optimal Control. Vol. 1 and 2. 4th edition, 2012.</li> <li>4. Algorithms for Reinforcement Learning, Csaba Szepesvári, Morgan &amp; Claypool, 2009.</li> <li>5. Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems, Sébastien Bubeck and Nicolò Cesa-Bianchi, Foundations and Trends in Machine Learning, Volume 5, Number 1, 2012.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Statistical Machine Learning</b> <b>(2-0-2-6)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Probability Models, Stochastic Process and its Applications, Math for Data Science I and II
3	<b>Course content</b>	Bayesian Decision Making and Bayes Classifier, Parametric Estimation of Densities, General Linear Models, EM Algorithm, Discriminative Learning based Models, Dimensionality Reduction Techniques, Empirical risk minimization, Ensemble Methods - Bagging, Boosting, Clustering
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. R.O.Duda, P.E.Hart and D.G.Stork, Pattern Classification, John Wiley, 2001.</li> <li>2. C.M.Bishop, Pattern Recognition and Machine Learning, Springer, 2006.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Runtime Verification</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Not applicable
3	<b>Course content</b>	<ol style="list-style-type: none"> <li>1. Overview of Runtime Verification, and its comparison with other Formal Verification approaches.</li> <li>2. Fundamentals: Propositional and First-Order Logic, Temporal Logics (Linear and Metric)</li> <li>3. Propositional LTL and its variants: specification of properties, runtime verification strategies, expressibility, and monitorability.</li> <li>4. First Order LTL and its variants: specification of properties, runtime verification strategies, expressibility, and monitorability.</li> <li>5. Discussion of various state-of-the-art tools and case studies.</li> </ol>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. K. Havelund, D. Peled, "Runtime Verification: From Propositional to First Order Temporal Logic", Tutorial at International Conference on Runtime Verification, 2018.</li> <li>2. Ezio Bartocci, Yliès Falcone. "Lectures on Runtime Verification". Springer, 2018. ISBN: 978-3-319-75632-5</li> <li>3. Michael Huth, Mark Ryan, "Logic in Computer Science: Modelling and Reasoning about Systems", Cambridge University Press, 2004. ISBN: 978-0521543101</li> <li>4. Research publications on Runtime Verification.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Systems Bootcamp for ML</b> <b>(1-0-2-4)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to basics of computer programming.
3	<b>Course content</b>	<p>Systems Programming lab focuses on programming principles and skills for building systems software. It also focuses on how to build and optimize AI/ML models from a systems perspective.</p> <ul style="list-style-type: none"> <li>● Introduction to UNIX, shell programming, version control and management (SVN/Git)</li> <li>● Introduction to Libraries (Boost, STL)</li> <li>● Introduction to Profiling (Performance analysis at system level and application level), ML perf</li> <li>● Introduction to various DL frameworks and DL Inference engines, Hardware backends</li> <li>● Database fundamentals and programming</li> <li>● GPU Programming (CUDA, OpenCL)/</li> <li>● Automatic Code Generation - TVM Stack</li> <li>● Mobile-edge cloud computing (Computational offloading decisions)</li> <li>● Programming assignments/projects will be given related to the above topics.</li> </ul>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Unix Man Pages for all unix tools, <b>Advanced Bash Scripting Guide from the Linux Documentation Project</b> (<a href="http://www.tldp.org">www.tldp.org</a>)</li> <li>2. Loeliger, Jon and McCullough, Matthew, <b>Version Control with Git: Powerful tools and techniques for collaborative software development</b>, O'Reilly, 2012.</li> <li>3. Sommerville, Ian. <b>Engineering Software Products: An Introduction to Modern Software Engineering</b>, Pearson, May 2019.</li> <li>4. Rodriguez Andres : <b>Deep Learning Systems: Algorithms, Compilers, and Processors</b> for Large- scale Production, Morgan and Claypool publishers, 2020</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Topics in Design and Analysis of Algorithms</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Discrete Mathematics, Design and Analysis of algorithms, Data structures and Algorithms.
3	<b>Course content</b>	<p><b>Module 1:</b> <u>Iterated Improvement Paradigms</u>- Computational and Algorithmic Thinking, Matching Algorithms, Flow Algorithms (16 hours).</p> <p><b>Module 2:</b> <u>Approximation Algorithms</u>- Greedy Approximation, Local Search, Linear Programming, Duality Techniques (16 hours)</p> <p><b>Module 3:</b> <u>Randomized Algorithms</u>- Monte Carlo and Las Vegas types, Randomized Attrition, Randomized Incremental Design, Sampling, Chernoff type bounds and High Confidence Analysis, Abundance of witness for Monte Carlo algorithms, Number theoretic Algorithms (16 hours).</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. [OA] James B. Orlin, Ravindra K. Ahuja, and Thomas L. Magnanti, "Network Flows", Prentice Hall, 1993.</li> <li>2. [WS] David P. Williamson and David B. Shmoys, "The Design of Approximation Algorithms", Cambridge University Press, 2011.</li> <li>3. [MR] Rajeev Motwani and Prabhakar Raghavan, "Randomized Algorithms", Cambridge University Press, 1995.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Advanced Algorithms</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Discrete Mathematics, Design and Analysis of algorithms, Data structures and Algorithms
3	<b>Course content</b>	<p><b>Module 1:</b> <u>Iterated Improvement Paradigms</u>- Computational and Algorithmic Thinking, Matching Algorithms, Flow Algorithms (16 hours).</p> <p><b>Module 2:</b> <u>Approximation Algorithms</u>- Greedy Approximation, Local Search, Linear Programming, Duality Techniques (16 hours)</p> <p><b>Module 3:</b> <u>Randomized Algorithms</u>- Monte Carlo and Las Vegas types, Randomized Attrition, Randomized Incremental Design, Sampling, Chernoff type bounds and High Confidence Analysis, Abundance of witness for Monte Carlo algorithms, Number theoretic Algorithms (16 hours).</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. [OA] James B. Orlin, Ravindra K. Ahuja, and Thomas L. Magnanti, "Network Flows", Prentice Hall, 1993.</li> <li>2. [WS] David P. Williamson and David B. Shmoys, "The Design of Approximation Algorithms", Cambridge University Press, 2011.</li> <li>3. [MR] Rajeev Motwani and Prabhakar Raghavan, "Randomized Algorithms", Cambridge University Press, 1995.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Topics in Graph Theory</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Discrete Structures
3	<b>Course content</b>	Recap: fundamental concepts. Topics in factors, covering and packing, cuts, connectivity, coloring, planarity, perfect graphs, Ramsey theory, and random graphs.
4	<b>Texts/References</b>	<p><b>Textbook:</b></p> <ol style="list-style-type: none"> <li>1. Introduction to Graph Theory (2nd Edition), Douglas B. West. Prentice Hall.</li> </ol> <p><b>References:</b></p> <ol style="list-style-type: none"> <li>1. Algorithmic Graph Theory and Perfect Graphs (2nd Edition), Martin Charles Golumbic. Elsevier.</li> <li>2. Graph Theory (Graduate Texts in Mathematics, 5th Edition), Reinhard Diestel. Springer.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Topics in Parameterized Algorithms and Complexity</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Data Structures and Algorithms, Design and Analysis of Algorithms
3	<b>Course content</b>	Introduction. Kernelization, Bounded Search Trees, Iterative Compression, Treewidth, Advanced kernelization algorithms. Lower bounds: Fixed-parameter intractability, lower bounds based on ETH, lower bounds for kernelization. Parameterized Algorithms, Kernelization, and Complexity of Graph Modification Problems.
4	<b>Texts/References</b>	<p><b>Textbook:</b></p> <ol style="list-style-type: none"> <li>1. Parameterized Algorithms, Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Daniel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Sourabh. Springer. 2015.</li> </ol> <p><b>Reference:</b></p> <ol style="list-style-type: none"> <li>1. Parameterized Complexity, R. G. Downey, and M. R. Fellows. Springer Science and Business Media. 2012.</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Power Aware Computing (3-0-2-8)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Computer Architecture, Operating Systems
3	<b>Course content</b>	Introduction to Power and Energy, Power consumption modeling and estimation, Dynamic power management and DVFS, Leakage reduction techniques, circuit-level and Micro-architecture techniques, Power states and ACPI support, Memory/cache power optimizations. Software level techniques, GPU power modeling and optimizations.
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. S. Kaxiras, M. Martonosi, Computer Architecture Techniques for Power-Efficiency, Synthesis Lectures on Computer Architecture. Morgan &amp;C laypool publishers.</li> <li>2. Siva G. Narendra, Anantha Chandrakasan</li> <li>3. P. Leakage in Nanometer CMOS Technologies, Series on Integrated Circuits and Systems</li> <li>4. Rakesh Chadha, J Bhasker an ASIC Low Power Primer: Analysis, Techniques and Specification.</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Dataflow Processor Architecture (Guided Study) (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Computer Architecture
3	<b>Course content</b>	The philosophy of dataflow; static and dynamic approaches; contrast with conventional out-of-order pipelines; understanding different granularities of operations, appreciating the performance and power possibilities; understanding the caveats; studying particular example architectures; analyzing the fundamental concept in the light of modern trends in the semiconductor industry; analyzing the fundamental concept in the context of particular application classes.
4	<b>Texts/References</b>	<p>Papers(list not exhaustive);</p> <ol style="list-style-type: none"> <li>1 Exploring the potential of heterogeneous von Neumann/Dataflow execution models, Nowatzki et al., ISCA2015.</li> <li>2 Dataflow Machine Architecture. AH Veen, ACM Computing Surveys. 1986.</li> <li>3 Dataflow Architecture: Are dataflow computers commercially viable?, Kavi et al., IEEE Potentials 1992.</li> <li>4 Synchronous Dataflow Architectures for Network processors, Carlstrom et al., IEEE Micro, 2004.</li> <li>5 Dataflow architectures, Culler. Annual review of Computer Science 1986.</li> <li>6 An architectural comparison of dataflow systems .Srin. dataflow Computing: Theory and Practice 1992.</li> <li>7 The Manchester Prototype Dataflow Computer. Gurd et al. ACM. 1985.</li> <li>8 Monsoon: an Explicit Token-store Architecture. Papadopolous et al. ACM SIGARCH 1990.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Advanced Data Structures and Algorithms Lab</b> (0-0-3-3)
2	<b>Pre-requisite courses(s)</b>	None
3	<b>Course content</b>	Concepts on advanced data structures and algorithms used in practical life. The lab course is to implement all these data structures and algorithms using any programming language, preferably one among C/C++, Python, and Java.
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. T. H. Cormen, C. E. Leiserson, R. L. Rivest and C. Stein, Introduction to algorithms, 2022. MIT press.</li> <li>2. Online lecture notes by Jeff Erickson [<a href="https://jeffe.cs.illinois.edu/teaching/algorithms/">https://jeffe.cs.illinois.edu/teaching/algorithms/</a> ]</li> <li>3. 3. Steven Skiena, The Algorithm Design Manual, 2000, Springer.</li> <li>4. 4. J. Kleinberg and E. Tardos, Algorithm Design, 2005, Pearson</li> <li>5. 5. Mark Weiss, Data structures and algorithm analysis in C++ (Java), 2014, Pearson</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Advanced Data Structures and Algorithms</b> (3-0-0-6)
2	<b>Pre-requisite courses(s)</b>	None
3	<b>Course content</b>	<p><b>Module 1:</b> Basics: asymptotic notations, recurrences, basic data structures</p> <p><b>Module 2:</b> Advanced data structures: heaps, priority queues, hash tables, data structures based on trees</p> <p><b>Module 3:</b> Design paradigms and complexity analysis: divide and conquer, dynamic programming, greedy algorithms, amortized analysis</p> <p><b>Module 4:</b> Advanced topics: graph algorithms, string algorithms, geometric algorithms, complexity lower bounds, coping up with hard problems.</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. T. H. Cormen, C. E. Leiserson, R. L. Rivest and C. Stein, Introduction to algorithms, 2022. MIT press.</li> <li>2. Online lecture notes by Jeff Erickson [<a href="https://jeffe.cs.illinois.edu/teaching/algorithms/">https://jeffe.cs.illinois.edu/teaching/algorithms/</a> ]</li> <li>3. 3. Steven Skiena, The Algorithm Design Manual, 2000, Springer.</li> <li>4. 4. J. Kleinberg and E. Tardos, Algorithm Design, 2005, Pearson</li> <li>5. 5. Mark Weiss, Data structures and algorithm analysis in C++ (Java), 2014, Pearson.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Advanced Software Development Laboratory</b> (1-0-4-6)
2	<b>Pre-requisite courses(s)</b>	None
3	<b>Course content</b>	<p>Editing: Vim/emacs,  Presentation: latex, beamer,  Build, Integration and Deployment, Version Control, and Documentation: make, GitHub and git, doxygen  Programming: HTML, CSS, Shell scripting, AWK, SED  Exploring features of IDE (e.g. eclipse, vscode) using a high-level programming language such as Java, Python, C++, Debugging (using gdb, using IDE).  Unix/Linux basics: shell, file system, permissions, process hierarchy, process monitoring, ssh, scp, rsync, grep, find, head, tail, tar, cut, sort, I/O redirection, pipes  Profiling tools: (e.g., gprof, prof, perf, valgrind)</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Article title: The Linux Documentation Project, URL: <a href="http://www.tldp.org">www.tldp.org</a>, Website title: The Linux Documentation Project, Date accessed: March 22, 2024.</li> <li>2. Article title: Python Tutorial, URL: <a href="http://docs.python.org/3/tutorial/index.html">http://docs.python.org/3/tutorial/index.html</a>, Website title: Python Documentation, Date accessed: March 22, 2024.</li> <li>3. Article title: Oracle Java Documentation, URL: <a href="http://docs.oracle.com/javase/tutorial/">http://docs.oracle.com/javase/tutorial/</a>, Website title: The Java Tutorials, Date accessed: March 22, 2024.</li> <li>4. Latex - A document preparation system, 2/e, by Leslie Lamport, Addison-Wesley, 1994.</li> <li>5. Article title: GNU Make, URL: <a href="https://www.gnu.org/software/make/manual/html_node/index.html#Top">https://www.gnu.org/software/make/manual/html_node/index.html#Top</a>, Website title: GNU Make, Date accessed: March 22, 2024.</li> <li>6. Article title: Book, URL: <a href="http://git-scm.com">Git - Book (git-scm.com)</a>, Website title: Git, Date accessed: March 22, 2024.</li> <li>7. Article title: GDB: The GNU Project Debugger, URL: <a href="http://sourceware.org">GDB Documentation (sourceware.org)</a>, Website title: GDB Documentation, Date accessed: March 22, 2024.</li> <li>8. Article title: Valgrind User Manual, URL: <a href="http://valgrind.org">Valgrind</a>, Website title: Valgrind, Date accessed: March 22, 2024.</li> <li>9. Article title: Overview, URL: <a href="https://www.doxygen.nl/manual/index.html">https://www.doxygen.nl/manual/index.html</a>, Website title: Doxygen, Date accessed: March 22, 2024.</li> <li>10. URL: <a href="http://www.gnu.org/software/gprof/">GNU gprof</a>, Website title: GNU gprof, Date accessed: March 22, 2024.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Combinatorics and Probability</b> (3-0-0-6)
2	<b>Pre-requisite courses(s)</b>	None
3	<b>Course content</b>	<p>Combinatorics:  <b>Principles of counting :</b>  Rule of sum, rule of product, permutations, combinations. Partition, modular arithmetic. Double counting, generating functions. Binomial coefficients, binomial theorem, multinomial theorem</p> <p><b>Probability theory:</b>  Probability axioms and laws, random variables, binomial distribution, Poisson, exponential and normal distributions, expectations and moments, joint distribution, conditional distribution , conditional expectations, convergence of random variables, law of large numbers, central limit theorem, Markov chains.</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. W. Feller, An Introduction to Probability Theory and its Applications, Vol.1, 3rd Edition, John Wiley.</li> <li>2. G. R. Grimmett and D. R. Stirzaker, Probability and Random Processes, 4th Edition, Oxford University Press, 2020.</li> <li>3. N. L., Biggs, Discrete Mathematics, 2nd Edition, Oxford University Press, 1989.</li> <li>4. Jiří Matoušek, Invitation to Discrete Mathematics, 2nd Edition, Oxford University Press, 2008</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Introduction to Programming – 1</b> <b>(3-0-2-4)</b>
2	<b>Pre-requisite courses(s)</b>	--
3	<b>Course content</b>	<p>This course provides an introduction to problem solving with computers using C  Topics covered will include:  <b>Utilization:</b> Developer fundamentals such as editor, integrated programming environment, Unix shell, modules, libraries.  <b>Programming features:</b> Machine representation, data types, arrays and records, objects, expressions, control statements, iteration, procedures, functions and recursion, Pointers, Structures and basic I/O. <b>Applications:</b> Sample problems in engineering, science, text processing, and numerical methods.</p>
4	<b>Texts/References</b>	<p>The C Programming Language Brian W Kernighan, Dennis M Ritchie, Prentice Hall India , 2nd edition, 1988  Programming with C (Second Edition) Byron Gottfried, Schaum's Outlines Series, Tata-Mcgraw Hill, 2011  How to Solve It by Computer, by G. Dromey, Prentice- Hall, Inc., Upper Saddle River, NJ, 1982. How to Solve _It (2nd ed.), by Polya, G., Doubleday and co, 1957.  Let Us C, by Yashwant Kanetkar, Allied Publishers, 1998.</p>

1	<b>Title of the course</b> (L-T-P-C)	<b>Introduction to High Performance Computing</b> <b>(0.5-0-0-1)</b>
2	<b>Pre-requisite courses(s)</b>	Nil
3	<b>Course content</b>	<p>Introduction to High Performance Computing (HPC): Need for HPC and supercomputing in scientific and engineering applications, definitions, serial and parallel computations</p> <p>Taxonomy of Computer Architectures: Single Instruction Stream Single Data Stream(SISD), Single Instruction Stream Multiple Data Stream (SIMD), Multiple Instruction Stream Single Data Stream (MISD), and Multiple Instruction Stream Multiple Data Stream (MIMD) architectures</p> <p>Parallel Computing on parallel Machines</p> <p>Parallel Computing on SIMD, MISD and MIMD architectures; vector and matrix operations</p> <p>Embarrassingly Parallel Applications</p> <p>Monte Carlo methods, fractals and images processing</p> <p>Other Applications</p> <p>Numerical algorithms and applications in science and engineering.</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. E. Aubanel, 'Elements of Parallel Computing', 1<sup>st</sup> Edition, Chapman and Hall/CRC, 2016.</li> <li>2. A. Grama, G. Karypis, V. Kumar and Gupta, 'An Introduction to Parallel Computing: Design and Analysis of Algorithms', 2<sup>nd</sup> Edition, Addison Wesley, Reading, 2003.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Data Structures and Algorithms</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Computer Programming
3	<b>Course content</b>	Introduction: data structures, abstract data types, analysis of algorithms. Creation and manipulation of data structures: arrays, lists, stacks, queues, trees, heaps, hash tables, balanced trees, tries, graphs. Algorithms for sorting and searching, order statistics, depth-first and breadth-first search, shortest paths and minimum spanning tree.
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Introduction to Algorithms, 3rd edition, by T. Cormen, C. Leiserson, R. Rivest, C. Stein, MIT Press and McGraw-Hill, 2009.</li> <li>2. Data structures and algorithms in C++, by Michael T. Goodrich, Roberto Tamassia, and David M. Mount, Wiley, 2004.</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Data Structures and Algorithms Laboratory (0-0-3-3)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Computer Programming (CS 102)
3	<b>Course content</b>	The laboratory course for CS 211 is based on creating and manipulating various data structures and implementation of algorithms.
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Introduction to Algorithms, 3rd edition, by T. Cormen, C. Leiserson, R. Rivest, C. Stein, MIT Press and McGraw-Hill, 2009.</li> <li>2. Data structures and algorithms in C++, by Michael T. Goodrich, Roberto Tamassia, and David M. Mount, Wiley, 2004.</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Automata Theory (3-1-0-8)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Discrete Structures
3	<b>Course content</b>	Finite state machines (DFA/NFA/epsilon NFAs), regular expressions. Properties of regular languages. Myhill-Nerode Theorem. Non-regularity. Push down automata. Properties of context-free languages. Turing machines: Turing hypothesis, Turing computability, Nondeterministic, multi tape and other versions of Turing machines. Church thesis, recursively enumerable sets and Turing computability. Universal Turing machines. Unsolvability, The halting problem, partial solvability, Turing enumerability, acceptability and decidability, unsolvable problems about Turing Machines. Post's correspondence problem.
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Introduction to Automata Theory, Languages and Computation, by John. E. Hopcroft, Rajeev Motwani, J. D. Ullman, 3rd edition. Pearson. 2013.</li> <li>2. Elements of the Theory of Computation, by H.R. Lewis and C. H. Papadimitrou, 2nd Edition. Prentice Hall Inc, 1998.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Discrete Structures</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	--
3	<b>Course content</b>	<p>There are four modules in the course:</p> <ol style="list-style-type: none"> <li><b>1. Proofs and structures</b> Introduction, propositions, predicates, examples of theorems and proofs, types of proof techniques, Axioms, Mathematical Induction, Well-ordering principle, Strong Induction, Sets, Russell's paradox, infinite sets, functions, Countable and uncountable sets, Cantor's diagonalization technique, Relations, Equivalence relations, partitions of a set.</li> <li><b>2. Counting and Combinatorics</b> Permutations, combinations, binomial theorem, pigeon hole principle, principles of inclusion and exclusion, double counting. Recurrence relations, solving recurrence relations.</li> <li><b>3. Elements of graph theory</b> Graph models, representations, connectivity, Euler and Hamiltonian paths, planar graphs, Trees and tree traversals.</li> <li><b>4. Introduction to abstract algebra and number theory</b> Semigroups, monoids, groups, homomorphisms, normal subgroups, congruence relations. Ceiling, floor functions, divisibility. Modular arithmetic, prime numbers, primality theorems.</li> </ol>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>5. Discrete Mathematics and its applications with Combinatorics and graph theory, 7th edition, by Kenneth H Rosen. Special Indian Edition published by McGraw-Hill Education, 2017.</li> <li>6. Introduction to Graph Theory, 2nd Edition, by Douglas B West. Eastern Economy Edition published by PHI Learning Pvt. Ltd, 2002.</li> <li>7. Discrete Mathematics, 2nd Edition, by Norman L Biggs. Indian Edition published by Oxford University Press, 2003.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Design and Analysis of Algorithms</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Computer Programming and Utilization, Discrete Structures, Data Structures and Algorithms , Data Structures and Algorithms Laboratory
3	<b>Course content</b>	<p>Syllabus is divided roughly into 8 modules; each module roughly takes two weeks.</p> <p><b>Module 1:</b> Introduction Examples and motivation. Asymptotic complexity: informal concepts, formal notation, examples</p> <p><b>Module 2:</b> Searching in list: binary search, Sorting: insertion sort, selection sort, merge sort, quicksort, stability and other issues.</p> <p><b>Module 3:</b> Divide and conquer: binary search, recurrence relations. nearest pair of points, merge sort, integer multiplication, matrix multiplication.</p> <p><b>Module 4:</b> Graphs: Motivation, BFS, DFS, DFS numbering and applications, directed acyclic graphs, directed acyclic graphs, Shortest paths: unweighted and weighted, Single source shortest paths: Dijkstra, Minimum cost spanning trees: Prim’s algorithm, Kruskal’s Algorithm</p> <p><b>Module 5:</b> Union-Find data structure, Priority queues, heaps. Heap sort. Dijkstra/Prims revisited using heaps, Search Trees: Introduction Traversals, insertions, deletions Balancing</p> <p><b>Module 6:</b> Greedy algorithms: Greedy: Interval scheduling, Proof strategies, Huffman coding.</p> <p><b>Module 7:</b> Dynamic Programming: weighted interval scheduling, memoization, edit distance, longest ascending subsequence. matrix multiplication, shortest paths: Bellman Ford, shortest paths: Floyd Warshall</p> <p><b>Module 8:</b> Intractability: NP completeness, reductions, examples, Misc topics.</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Algorithms, by Sanjoy Dasgupta, Christos Papadimitriou and Umesh Vazirani, McGraw Hill Education, 2006.</li> <li>2. Introduction to Algorithms, 3rd edition, by Cormen, Leiserson, Rivest and Stein, PHI Learning Pvt. Ltd., 2010.</li> <li>3. Algorithm Design, 1st edition, by Kleniberg and Tardos, Pearson, 2014.</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Artificial Intelligence (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	
3	<b>Course content</b>	<p>Search: Problem representation; State Space Search; A* Algorithm and its Properties; AO* search, Minimax and alpha- beta pruning, AI in games. Logic: Formal Systems; Notion of Proof, Decidability, Soundness, Consistency and Completeness; Predicate Calculus (PC), Resolution Refutation, Herbrand Interpretation, Prolog. Knowledge Representation: PC based Knowledge Representation, Intelligent Question Answering, Semantic Net, Frames, Script, Conceptual Dependency, Ontologies, Basics of Semantic Web. Learning: Learning from Examples, Decision Trees, Neural Nets, Hidden Markov Models, Reinforcement Learning, Learnability Theory. Uncertainty: Formal and Empirical approaches including Bayesian Theory, Fuzzy Logic, Non-monotonic Logic, Default Reasoning. Planning: Blocks World, STRIPS, Constraint Satisfaction, Basics of Probabilistic Planning.</p> <p>Advanced Topics: Introduction to topics like Computer ain</p>
4	<b>Texts/References</b>	<p>Text: Stuart J. Russel, Peter Norvig, Artificial Intelligence: A Modern Approach (3rd ed.). Upper Saddle River: Prentice Hall, 2010. Other references: N.J. Nilsson, Principles of Artificial Intelligence, Morgan Kaufmann, 1985. Malik Ghallab, Dana Nau, Paolo Traverso, Automated Planning: Theory &amp; Practice, The Morgan Kaufmann Series in Artificial Intelligence, 2004. Christopher Bishop, Pattern Recognition and Machine Learning, Springer, 2006. Mark Stefik, Introduction to Knowledge Systems, Morgan Kaufmann, 1995. E. Rich and K.Knight, Artificial Intelligence, Tata McGraw Hill, 1992.</p>

1	<b>Title of the course</b> (L-T-P-C)	<b>Software Systems Laboratory</b> <b>(1-3-0-8)</b>
2	<b>Pre-requisite courses(s)</b>	--
3	<b>Course content</b>	<p>Vim/emacs HTML, CSS</p> <ul style="list-style-type: none"> <li>• Report and presentation software: latex, beamer, drawing software (e.g. inkscape, xfig, open-office)</li> <li>• IDE (e.g. eclipse, netbeans), code reading, debugging Basic Java Java collections, interfaces</li> <li>• Java threads Java GUI Introduction to documentation: e.g. doxygen/javadocs</li> <li>• Version management: SVN/Git</li> <li>• Unix basics: shell, file system, permissions, process hierarchy, process monitoring, ssh, rsync</li> <li>• Unix tools: e.g. awk, sed, grep, find, head, tail, tar, cut, sort</li> <li>• Bash scripting: I/O redirection, pipes</li> <li>• Python programming</li> <li>• Makefile, libraries and linking</li> <li>• Graph plotting software (e.g., gnuplot)</li> <li>• Profiling tools (e.g., gprof, prof)</li> <li>• Optional topics (may be specific to individual students 30222 projects): intro to sockets, basic SQL for data storage, JDBC/pygresql</li> </ul> <p>A project would be included which touches upon many of the above topics, helping students see the connection across seemingly disparate topics. The project is also expected to be a significant load: 20-30 hours of work.</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Online tutorials for HTML/CSS, Inkscape, OODraw Unix Man Pages for all unix tools, Advanced Bash Scripting Guide from the Linux Documentation Project (<a href="http://www.tldp.org">www.tldp.org</a>).</li> <li>2. The Python Tutorial Online Book (<a href="http://docs.python.org/3/tutorial/index.html">http://docs.python.org/3/tutorial/index.html</a>).</li> <li>3. The Java Tutorials (<a href="http://docs.oracle.com/javase/tutorial/">http://docs.oracle.com/javase/tutorial/</a>).</li> <li>4. Latex - A document preparation system, 2/e, by Leslie Lamport, Addison-Wesley, 1994.</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Artificial Intelligence Lab (0-0-3-3)</b>
2	<b>Pre-requisite courses(s)</b>	
3	<b>Course content</b>	The lab will closely follow and aim to elucidate the lessons covered in the theory course CS344. Implementation and study of A*, Usage of Prolog Inferencing, Expert System Shells, Neural Net Platforms, Prediction and Sequence Labeling using HMMs, Simulation of Robot Navigation and such exercises are strongly recommended.
4	<b>Texts/References</b>	Stuart J. Russel, Peter Norvig, Artificial Intelligence: A Modern Approach (3rd ed.). Upper Saddle River: Prentice Hall, 2010. Other references: N.J. Nilsson, Principles of Artificial Intelligence, Morgan Kaufmann, 1985. Malik Ghallab, Dana Nau, Paolo Traverso, Automated Planning: Theory & Practice, The Morgan Kaufmann Series in Artificial Intelligence, 2004. Christopher Bishop, Pattern Recognition and Machine Learning, Springer, 2006. Mark Stefik, Introduction to Knowledge Systems, Morgan Kaufmann, 1995. E. Rich and K.Knight, Artificial Intelligence, Tata McGraw Hill, 1992.

1	<b>Title of the course (L-T-P-C)</b>	<b>Computer Architecture (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	--
3	<b>Course content</b>	<p>The Language of Bits, Assembly Language, LogicGates, Registers, and Memories, Processor Design, Principles of Pipelining, The Memory System, Multiprocessor Systems, I/O and Storage Devices.</p> <p>Each concept will be first taught on the basis of the fundamental driving principles. Following this, real world examples (e.g., ARM processors) will be used to emphasize the content.</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Computer Organization and Architecture, by SmrutiRanjan Sarangi, McGraw Higher Ed, 2017.</li> <li>2. Computer Architecture A Quantitative Approach, Sixth edition, by David Patterson and John L. Hennessy, Morgan Kaufmann, 2017.</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Data Bases and Information Systems (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	--
3	<b>Course content</b>	<p>Overview of data management systems. Relational model and query languages (relational algebra and calculus, SQL). Database design using the ER Model, ER Diagrams, UML Class Diagrams. Relational database design and normalization. Integrity and Security. Design and development of Web based information systems. Overview of storage structures and indexing, query processing and optimization, and transaction processing. Introduction to Big Data management concepts such as: distributed and scaledata storage, including distributed file systems, keyvalue stores, column stores and graph databases, replication and consistency, and concurrent data processing using the Map Reduce paradigm.</p> <p>Introduction to decision support and data analysis, data warehousing and data mining, and Information Retrieval.</p>
4	<b>Texts/References</b>	Database System Concepts, 6th edition, by AbrahamSilberschatz, Henry F. Korth and S. Sudarshan, McGraw Hill, 2010.

1	<b>Title of the course</b> (L-T-P-C)	<b>Operating Systems</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Computer Architecture
3	<b>Course content</b>	Process Management, Memory Management, Storage Management, Protection and Security, Virtual Machines, Distributed Systems
4	<b>Texts/References</b>	Avi Silberschatz, Peter Baer Galvin, Greg Gagne, ``Operating Systems Concepts" 9th edition. <i>Wiley</i> . Andrew S. Tanenbaum, Herbert Bos, ``Modern Operating Systems", 4th edition. <i>Pearson</i>

1	<b>Title of the course (L-T-P-C)</b>	<b>Introduction to Artificial Neural Networks (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Nil
3	<b>Course content</b>	<p>Background to ANN and PDP models; Basics of ANN including terminology, topology and learning laws; (4 lectures)</p> <p>Analysis of Feedforward neural networks (FFNN) including linear associative networks, perceptron network, multilayer perceptron, gradient descent methods and backpropagation learning; (8 lectures)</p> <p>Analysis of Feedback neural networks (FBNN) including Hopfield model, state transition diagram, stochastic networks, Boltzmann learning law; (8 lectures)</p> <p>Evolution of ANN architectures - from learning to deep learning; (1 lecture)</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. B Yegnanarayana, Artificial Neural Networks, Prentice Hall of India, New Delhi, 1999.</li> <li>2. David E Rumelhart, James L McClelland, and the PDP Research group, Eds, Parallel and Distributed Processing: Explorations in Microstructure of Cognition, Vol.1, Cambridge MA: MIT Press, 1986a</li> <li>3. James L McClelland, David E Rumelhart and the PDP Research group, Eds, Parallel and Distributed Processing: Explorations in Microstructure of Cognition, Vol.2, Cambridge MA: MIT Press, 1986b</li> <li>4. James L McClelland, David E Rumelhart and the PDP group, Eds, Explorations in Parallel and Distributed Processing: A Handbook of Models, Cambridge MA: MIT Press, 1989</li> <li>5. Simon Haykin, Neural Networks and Learning Machines, Pearson Education, 2011</li> <li>6. Ian Goodfellow, Yoshua Bengio and Aaron Courville, Deep learning, MIT Press, 2017.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Computer Architecture Laboratory</b> <b>(0-0-3-3)</b>
2	<b>Pre-requisite courses(s)</b>	--
3	<b>Course content</b>	The lab will closely follow the theory course. The idea isto have the students develop a software model of a simple processor, capturing both functionality and timing aspects. Theywill implement modules as the concepts aretaught in class.
4	<b>Texts/References</b>	Nil

1	<b>Title of the course (L-T-P-C)</b>	<b>Data Bases and Information Systems Laboratory (0-0-3-3)</b>
2	<b>Pre-requisite courses(s)</b>	--
3	<b>Course content</b>	Use of database systems supporting interactive SQL. Two-tier client-server applications using JDBC or ODBC, Three-tier web applications using Java servlets/JDBC orequivalent. Design of applications and user interfaces using these systems. Data analysis tools. Laboratory project involving building data backed applications withWeb or mobile app frontends.
4	<b>Texts/References</b>	Abraham Silberschatz, Henry F. Korth and S. Sudarshan, Database System Concepts 6th Ed, McGrawHill, 2010.

1	<b>Title of the course</b> (L-T-P-C)	<b>Operating Systems Laboratory</b> <b>(0-0-3-3)</b>
2	<b>Pre-requisite courses(s)</b>	<b>Computer Architecture</b>
3	<b>Course content</b>	Laboratory Assignments related to the topics covered in the theory course: Process Management, Memory Management, Storage Management, Protection and Security, Virtual Machines, Distributed Systems
4	<b>Texts/References</b>	1. <b>Avi Silberschatz, Peter Baer Galvin, Greg Gagne</b> , ``Operating Systems Concepts" 9th edition. Wiley. 2. <b>Andrew S. Tanenbaum, Herbert Bos</b> , ``Modern Operating Systems", 4th edition. Pearson.

1	<b>Title of the course (L-T-P-C)</b>	<b>Computer Networks Laboratory (0-0-3-3)</b>
2	<b>Pre-requisite courses(s)</b>	Nil
3	<b>Course content</b>	Experiments to support study of the Internet protocol stack: (a) Experimental study of application protocols such as HTTP, FTP, SMTP, using network packet sniffers and analyzers such as Ethereal. Small exercises in socket programming in C/C++/Java. (b) Experiments with packet sniffers to study the TCP protocol. Using OS (netstat, etc) tools to understand TCP protocol FSM, retransmission timer behavior, congestion control behaviour. (c) Introduction to ns2 (network simulator) - small simulation exercises to study TCP behavior under different scenarios. (d) Setting up a small IP network - configure interfaces, IP addresses and routing protocols to set up a small IP network. Study dynamic behaviour using packet sniffers (e) Experiments with ns2 to study behaviour (especially performance of) link layer protocols such as Ethernet and 802.11 wireless LAN.
4	<b>Texts/References</b>	Nil

1	<b>Title of the course (L-T-P-C)</b>	<b>Compilers Lab (0-0-3-3)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Data Structures and Algorithms, Computer Architecture, Automata Theory, and a programming language such as C/C++/Java.
3	<b>Course content</b>	Design and implementation of a scanner using scanner generator. Design and implementation of a parser using parser generator. Symbol table generation, Semantic actions for expressions, control structures, and functions. Implementing liveness analysis and applying it to register allocation.
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Alfred V. Aho, Monica S. Lam, Ravi Sethi and Jeffrey D.Ullman: Compilers: Principles, Techniques, and Tools, 2/E, AddisonWesley 2007.</li> <li>2. Andrew Appel: Modern Compiler Implementation in C/ML/Java, Cambridge University Press, 2004</li> <li>3. Dick Grune, Henri E. Bal, Cerial J.H. Jacobs and Koen G. Langendoen: Modern Compiler Design, John Wiley &amp; Sons, Inc. 2000.</li> <li>4. Michael L. Scott: Programming Language Pragmatics, Morgan Kaufman Publishers, 2006.</li> <li>5. Fisher and LeBlanc: Crafting a Compiler in C.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Compilers</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Data Structures and Algorithms, Computer Architecture, Automata Theory
3	<b>Course content</b>	The compiled and interpreted execution models. Lexical analysis and parsing using lex and yacc. LR parsers, Scope and visibility analysis. Data layout and lifetime management of data. Runtime environment. Dynamic memory allocation and Garbage collection. Translation of expressions, control structures, and functions. Code generation and introduction to optimizations (local and global). Lattice Theory, Optimizations- dataflow, control flow, reaching definition, liveness analysis, code transformation-tiling, fusion.
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Alfred V. Aho, Monica S. Lam, Ravi Sethi and Jeffrey D.Ullman: Compilers: Principles, Techniques, and Tools, 2/E, AddisonWesley 2007.</li> <li>2. Andrew Appel: Modern Compiler Implementation in C/ML/Java, Cambridge University Press, 2004</li> <li>3. Dick Grune, Henri E. Bal, Cerial J.H. Jacobs and Koen G. Langendoen: Modern Compiler Design, John Wiley &amp; Sons, Inc. 2000. Michael L. Scott: Programming Language Pragmatics, Morgan Kaufman Publishers, 2006.</li> <li>4. Fisher and LeBlanc: Crafting a Compiler inC.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Programming Techniques</b> <b>(3-0-0-3)</b>
2	<b>Pre-requisite courses(s)</b>	A first-level course in any programming language.
3	<b>Course content</b>	<p><b>The programming Techniques</b> course starts with refresher topics on programming language constructs and data structures. Then, the discussion switches to data structures and algorithms for problem solving. The course is meant for students who are familiar with a programming language (preferably C++ and Python) and those who desire to improve techniques of problem solving and programming.</p> <ol style="list-style-type: none"> <li>1. Programming Refresher</li> <li>2. Data structures refresher, ADTs <ol style="list-style-type: none"> <li>1. Arrays, lists, multilists, queues, and stacks.</li> </ol> </li> <li>3. Non-linear data structures <ol style="list-style-type: none"> <li>1. Trees, Graphs, Hash Tables, Tries, Suffix trees</li> </ol> </li> <li>4. Algorithms <ol style="list-style-type: none"> <li>1. Basic analysis, complexity</li> <li>2. Design techniques</li> </ol> </li> <li>5. Some common problems and program design strategies</li> </ol>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Course Notes</li> <li>2. How to Solve it by computers: Dromey, Pearson, 2006</li> <li>3. The C Programming Language: Kernighan and Ritchie, Second Edition, Pearson, 2015</li> <li>4. Data Structures and program design: Kruse, Pearson, 2006</li> <li>5. Introduction to the Design and Analysis of Algorithms, Anany Levitin, third edition, 2017</li> <li>6. Programming Pearls, Jon Bentley, 2002.</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Computer Networks (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Nil
3	<b>Course content</b>	Design of Computer Networking protocols at all layers: transmission media, data link protocols, media access control, routing and congestion control, admission control, traffic shaping and policing, Internet working (IP) and transport layer protocols (TCP). Performance analysis of networks.
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Data and Computer Communications, 6th edition, by W. Stallings, Prentice Hall, 2000.</li> <li>2. Computer Networks, 4th edition, by A. S. Tannenbaum, Prentice Hall, 2003.</li> <li>3. Data Communications, Computer Networks and Open Systems, 4th edition, by F. Halsall, Addison-Wesley, 1996.</li> <li>4. High Performance Communication Networks, by Walrand and Varaiya, Morgan Kaufman, 1996.</li> <li>5. Internet working with TCP/IP: Principles, Protocols, Architecture, 3rd edition, by D. E. Comer, Prentice Hall, 1996.</li> <li>6. TCP/IP Illustrated Vol. I, by W. R. Stevens, Addison Wesley, 1994.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Software Engineering</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	
3	<b>Course content</b>	<p><b>Introduction</b> What is Software Engineering.</p> <p><b>Software Development Lifecycle</b> Requirements analysis, software design, coding, testing, maintenance, etc.</p> <p><b>Software life-cycle models</b> Waterfall model, prototyping, interactive enhancement, spiral model. Role of Management in software development. Role of metrics and measurement.</p> <p><b>Software Requirement Specification</b> Problem analysis, requirement specification, validation, metrics, monitoring and control.</p> <p><b>System Design</b> Problem partitioning, abstraction, top-down and bottom-up design, Structured approach. Functional versus object-oriented approach, design specification and verification metrics, monitoring and control. Software Architecture</p> <p><b>Coding</b> Top-down and bottom-up, structured programming, information hiding, programming style, and internal documentation. Verification, Metrics, monitoring and control.</p> <p><b>Testing</b> Levels of testing functional testing, structural testing, test plane, test cases specification, reliability assessment.</p> <p><b>Software Project Management</b> Cost estimation, Project scheduling, Staffing, Software configuration management, Quality assurance, Project Monitoring, Risk management, etc. including tools for software development to release, supporting the whole life cycle.</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Software Engineering: A Practioner’s approach, R.S. Pressman, McGraw Hill, 8th edition.</li> <li>2. Introduction to Software Engineering, Pankaj Jalote, Narosha Publishing</li> <li>3. The Unified Software Development Process, I. Jacobson, G. Booch, J. Rumbaugh, Pearson Education</li> <li>4. Software Architecture in Practice, L. Bass, P. Clements, R. Kazmann, 3rd ed., Addison Wesley</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Distributed Systems</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Operating Systems, Data Structures and Algorithms, Programming in C++
3	<b>Course content</b>	<ul style="list-style-type: none"> <li>• Introduction to distributed systems, Message Passing, Leader Election, Distributed Models, Causality and Logical Time</li> <li>• Logical Time, Global State &amp; Snapshot and Distributed Mutual Exclusion-Non- Token and Quorum based approaches</li> <li>• Distributed Mutual Exclusion-Token based approaches, Consensus &amp; Agreement, Checkpointing &amp; Rollback Recovery</li> <li>• Deadlock Detection, DSM and Distributed MST</li> <li>• Termination Detection, Message Ordering &amp; Group Communication, Fault Tolerance and Self-Stabilization, Gossip Style communication, chord, pastry</li> <li>• Concurrency and Replication Control, RPCs, Transactions</li> <li>• Distributed Randomized Algorithms, DHT and P2P Computing</li> <li>• Case Studies: GFS, HDFS, Map Reduce and Spark</li> </ul>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Distributed Computing: Principles, Algorithms, and Systems- Ajay D. Kshemkalyani and Mukesh Singhal</li> <li>2. Distributed Computing: Fundamentals, Simulations and Advanced Topics- Hagit Attiya and Jennifer Welch</li> <li>3. Distributed Algorithms-Nancy Lynch</li> <li>4. Elements of Distributed Computing-Vijay K. Garg</li> <li>5. Advanced Concepts in Operating Systems-Mukesh Singhal, Niranjan G. Shivaratri</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Graph Theory and Combinatorics (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Discrete Structures
3	<b>Course content</b>	Fundamentals of graph theory. Topics include: connectivity, planarity, perfect graphs, coloring, matchings and extremal problems. Basic concepts in combinatorics. Topics include: counting techniques, inclusion-exclusion principles, permutations, combinations and pigeon-hole principle.
4	<b>Texts/References</b>	<p>“An Introduction to Quantum Field Theory”, Michael Peskin and Daniel Schroeder (Addison Wesley)</p> <p>“Introduction to Quantum Field Theory”, A. Zee</p> <p>“Quantum Field Theory”, Lewis H. Ryder</p> <p>“Quantum Field Theory and Critical Phenomena”, by Jean Zinn-Justin.</p> <p>“Quantum field Theory for the Gifted Amateur”, T. Lancaster and Stephen J. Blundell</p> <p>NPTEL lectures in Quantum Field Theory (<a href="https://nptel.ac.in/courses/115106065/">https://nptel.ac.in/courses/115106065/</a>)</p>

1	<b>Title of the course</b> (L-T-P-C)	<b>Parallel Computing</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to C, C++ or Fortran programming
3	<b>Course content</b>	<p>Need for High Performance Computing (HPC) and applications.</p> <p>Sequential Computing model, Algorithms and their complexity.</p> <p>Taxonomy of computer architectures – SISD, SIMD (e.g. array processors), MISD (pipelined processing, vector processors), and MIMD (shared memory and distributed memory multiprocessors, computing clusters); dataflow computing; hardware accelerators (GPUs); interconnection networks (bus, loop, mesh and hypercube); Memory hierarchy; Case Studies.</p> <p>Implications of computer architectures to algorithm design, synchronous processing, single program multiple data (SPMD) and multiple program multiple data (MPMD) processing; functional and data parallelism; memory hierarchies.</p> <p>Performance evaluation: communication and computing costs, speedup, efficiency, Amdahl’s law, parallel scalability.</p> <p>Parallel algorithm design and case studies: numerical algorithms (linear algebra, matrix- vector and matrix-matrix multiplications, finite difference method and PDEs, Monte Carlo method), and non-numerical algorithms (search, sorting, simple tree and graph algorithms)</p> <p>Parallel programming platforms, OpenMP and MPI programming, GPU programming.</p> <p>Programing Assignments:</p> <ol style="list-style-type: none"> <li>1. Parallel computing lab environment (system architecture, log on, hello world</li> <li>2. Editors, job submission, optimization techniques for serial code.</li> <li>3. MPI and simple program(s)</li> <li>4. MPI and matrix-matrix multiplication</li> <li>5. OpenMP and matrix-matrix multiplication OpenMP</li> <li>6. Introduction to GPU programming – matrix-matrix multiplication.</li> </ol>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Ananth Grama, Anshul Gupta, George Karypis, Vipin Kumar: Introduction to Parallel Computing, Addison Wesley 2003</li> <li>2. Eric Aubanel, Elements of Parallel Computing, CRC Press, 2017.</li> <li>3. <a href="https://computing.llnl.gov/tutorials/mpi/">https://computing.llnl.gov/tutorials/mpi/</a></li> <li>4. <a href="https://computing.llnl.gov/tutorials/open_MP/">https://computing.llnl.gov/tutorials/open MP/</a></li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Principles of Programming Languages (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Knowledge of a programming language, preferably multiple multi-paradigm programming languages. Mathematical and analytical inclination.
3	<b>Course content</b>	<p><b>Foundations of Programming Languages</b> is an undergraduate introductory course that covers formal aspects of computations and programming language constructs. While the course discusses formal aspects, it also draws from different programming languages like Haskell, ML, Clojure, Java, Scala and C++. Some of the topics that will be covered are:</p> <ol style="list-style-type: none"> <li>1. Types and introduction to type theory</li> <li>2. Recursive Types</li> <li>3. Untyped and typed lambda calculus</li> <li>4. Types and Subtypes and polymorphism</li> <li>5. Object Calculus</li> <li>6. Concurrent Programming</li> <li>7. Reflective Programming</li> </ol>
4	<b>Texts/References</b>	<p>Text:</p> <ol style="list-style-type: none"> <li>1. Essentials of Programming Languages, Friedman, Wand and Haynes, PHI, 2012.</li> <li>2. Foundations of Object-Oriented Languages, Kim Bruce, PHI 2004. Papers from literature</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Elements of Programming</b> (2-0-2-6)
2	<b>Pre-requisite courses(s)</b>	None
3	<b>Course content</b>	<p><b>Elements of Programming</b> is a first-level course in programming. Starting with the basics of programming, it culminates in fairly complex programs covering program design, structure and techniques along the way. This course uses Python programming language to learn the principles. Topics covered include:</p> <ol style="list-style-type: none"> <li>1. Program structure</li> <li>2. Types and operations on types</li> <li>3. Abstractions in programs</li> <li>4. Object-Oriented Programming</li> <li>5. Functional Programming</li> <li>6. Exceptions</li> <li>7. Program performance</li> <li>8. Problem Solving techniques – data structures and algorithms</li> <li>9. Programming for data analytics</li> </ol>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Required Text: Introduction to Computation and Programming Using Python, 2<sup>nd</sup> Edition, by John V. Guttag, Prentice Hall of India, 2016</li> <li>2. Reference: 1. Think Like a Programmer, by V. Anton Spraul, No Starch Press, 2012</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Introduction to Blockchains (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	It is expected that students should have a very good programming background preferably in Java/JavaScript and Python.
3	<b>Course content</b>	<p>Basics - What is Blockchain, Centralized vs Decentralized vs Distributed, Blockchains and Public Ledgers,  Fundamentals - Cryptographically secure hash functions, Merkle Trees, Bitcoin Concept  Architecture - Blockchain 2.0, smart contracts, block structure, notion of distributed consensus, challenge response to Permission-less Consensus, economics behind blockchain consensus  Consensus in Blockchain - Distributed Consensus, Proof of Work, Miners in the context of Bitcoin  Permissioned Blockchain - Basics, RAFT Consensus, Byzantine General Problem, Practical Byzantine Fault Tolerance  Hyperledger Fabric - Introduction, Transaction flow, Membership and Identity Management, Hyperledger Composer"  Ethereum Framework - Installation and subsequent execution of the use cases.  Blockchain in Financial Service - Payments and Secure Trading, Compliance and Mortgage, Financial Trade  Blockchain in Supply Chain  Blockchain in Government Use Cases - Digital Identity"  Mini Project implementation using Ethereum/Hyperledger framework on use cases related to financial, supply chain, and government sectors</p>
4	<b>Texts/References</b>	<p>Arvind Narayanan, Joseph Bonneau, Edward Felten, Andrew Miller, Steven Goldfeder, "Bitcoin and Cryptocurrency Technologies – A Comprehensive Introduction", Princeton University Press, 2016.</p> <p>Research papers as well as several related online educational content will also be referred to for learning some of the afore-mentioned topics.</p>

<b>1</b>	<b>Title of the course (L-T-P-C)</b>	<b>Logic for Computer Science (3-0-0-6)</b>
<b>2</b>	<b>Pre-requisite courses(s)</b>	Discrete Mathematics, Theory of computation.
<b>3</b>	<b>Course content</b>	<ol style="list-style-type: none"> <li>1. Module 1 :Propositional Logic: Syntax, Semantics, Normal Forms, Boolean Functions.</li> <li>2. Module 2: Computational complexity of Satisfiability P vs NP, SAT: hardest among NP.</li> <li>3. Module 3: Syntactic SAT solvers : Resolution, Tableaux.</li> <li>4. Module 4:proof Systems: Semantic entailment, Compactness, Soundness Completeness, Natural Deduction, Gentzen Sequent Calculus, Hilbert System.</li> <li>5. Module 5: Predicate Logic. Randomized SAT solvers. Programming assignments: using SAT/SMT solver z3.</li> </ol>
<b>4</b>	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Logic in Computer Science, Michael Huth and Mark Ryan, Cambridge University Press.</li> <li>2. SAT/SMT by example, Dennis Yurichev.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Mathematics for Data Science</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to basic concepts in calculus and linear algebra
3	<b>Course content</b>	<p>Introduction to Data science and Motivation for the course.</p> <p>Review of calculus, naïve set theory, notion of limits, ordering, series and its convergence. Introduction to Linear Algebra in Data science, notion of vector space, dimension and rank, algorithms for solving linear equations, importance of norms and notion of convergence, matrix decompositions and its use.</p> <p>Importance of optimization in data science: Birds view of Linear Regression, Multivariate Regression, Logistic Regression etc.</p> <p>Convex Optimization: Convex sets, convex functions, duality theory, different types of optimization problems, Introduction to linear program.</p> <p>Algorithms: Central of gravity method, Gradient descent methods, Nesterov acceleration, mirror descent/ Nesterov dual averaging, stochastic gradient methods, Rmsprop, SIGNSGD, ADAM algorithm etc.</p> <p>Non-convex optimization: Demonstration of convex methods on non-convex problems; merits and disadvantages.</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. C. Bishop, "Pattern Recognition and Machine Learning," Springer, 2006.</li> <li>2. Cambridge university press, 2018 (reprint). for Machine Learning," Now publisher, 2017.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Deep Learning</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Calculus, Linear Algebra, Probability, Random Processes, Ability to code in Python
3	<b>Course content</b>	<p><b>Introductory Concepts of DNN</b></p> <ol style="list-style-type: none"> <li>1. Linear regression, logistic regression –penultimate layers of a neural network</li> <li>2. Dealing with nonlinearity – Kernal Trick</li> <li>3. Data-driven kernel learning using NNs</li> </ol> <p><b>DNN Training</b></p> <ol style="list-style-type: none"> <li>1. Issues in training practical deep networks, Vanishing/Exploding gradients</li> <li>2. Regularization for Deep Learning – Early stopping, weight regularization, activity regularization, dropout</li> <li>3. Optimization methods for training deep networks – Stochastic gradient descent, rmsprop, adam</li> <li>4. Convolutional Neural networks</li> </ol> <p><b>Sequence Modeling</b></p> <ol style="list-style-type: none"> <li>1. Recurrent neural network</li> <li>2. LSTMs ans BLSTMs</li> </ol> <p><b>Unsupervised Learning</b></p> <ol style="list-style-type: none"> <li>1. Autoencoder</li> <li>2. Variational autoencoder</li> <li>3. Generative adversial networks (GANs)</li> <li>4. Representation learning and feature extraction.</li> </ol>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Ian Goodfellow and Yoshua Bengio and Aaron Courville, “Deep Learning,” MIT Press</li> <li>2. Bishop, C. M. Neural Networks for Pattern Recognition. Oxford University Press. 1995</li> <li>3. B Yegnanarayana, “Artificial Neural Networks,” PHI.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Computer Graphics</b> <b>(3-0-2-8)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Computer programming and basic linear algebra.
3	<b>Course content</b>	Introduction to the course, Rasterization Basics, Drawing in OpenGL, Clipping 2D Transformations, 3D Transformations, Viewing Transformations, The Modeling-Viewing Pipeline, Visibility, Hierarchical Modelling, Shading, Texture Cubic Splines, Bezier Splines, B-Splines Principles of Animation, Interpolation for Animation, Modelling Surfaces
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Fundamentals of Computer Graphics (Third Edition), Peter Shirley, Steve Marschner and others, A K Peters/CRC Press (2009)</li> <li>2. Interactive Computer Graphics - A Top- Down Approach Using OpenGL (6/e), Edward Angel</li> <li>3. Computer Graphics using OpenGL (3/e), F. S.Hill Jr. and S. M. Kelley</li> <li>4. Computer Graphics with OpenGL (3/e), D. D. Hearn and M. P. Baker</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Introduction to Logic (3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to Discrete Mathematics
3	<b>Course content</b>	<ol style="list-style-type: none"> <li>1. Propositional logic: Natural deduction, Semantics of propositional logic, Soundness of propositional logic, Completeness of propositional logic Horn clauses and satisfiability.</li> <li>2. Predicate logic: Natural deduction rules, decidability of predicate logic, Expressiveness of predicate logic</li> <li>3. Program correctness: Hoare triples, Partial and total correctness, Proof calculus for partial correctness, Proof calculus for total correctness</li> <li>4. Other Applications such as Logic in databases, Logic programming, Puzzle solving</li> <li>5 Practice with Verification tools.</li> </ol>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Logic in Computer Science. Huth and Ryan. Cambridge University Press, 2004</li> <li>2. A Mathematical Introduction to Logic. Herbert D Enderton. Harcourt Academic Press.</li> <li>3. First-Order Logic and Automated Theorem Proving by Melvin Fitting.</li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Cloud Software Development</b> <b>(1.5-0-0-3)</b>
2	<b>Pre-requisite courses(s)</b>	<b>Desirable:</b> Exposure on Operating System, Database, Cloud Programming language (Java, .Net, NodeJS, HTML/CSS, etc.)
3	<b>Course content</b>	<p><b>Module 1 - Introduction to Cloud Computing Landscape</b></p> <ul style="list-style-type: none"> <li>● Understand how industries rely on the cloud computing global infrastructure, Identify the applications and use cases</li> <li>● Identify the principles and characteristics of Cloud Computing - IaaS, PaaS, SaaS</li> <li>● Validate the different patterns of cloud computing adoption including public cloud services, private and hybrid approaches</li> <li>● Identify common challenges associated with the adoption of cloud computing solutions and associated myths</li> <li>● Compare and contrast with on-premise/traditional versus cloud</li> <li>● Understand in-country data regulations, data sovereignty considerations</li> </ul> <p><b>Module 2 - Cloud Computing Technology</b></p> <ul style="list-style-type: none"> <li>● Understand Virtualization Concepts - data, compute, network, operating system, HCI</li> <li>● Understand Cloud Infrastructure -Backup, Restore, Migration, DC/DR, HA use cases</li> <li>● Understand Programming concepts Cloud-native apps, Serverless, Containers</li> <li>● Learn Containers– Kubernetes, Docker, containers</li> </ul> <p><b>Module 3 - Using Managed Cloud Services</b></p> <ul style="list-style-type: none"> <li>● Learn 12-factor Application Architecture, api, Microservices, databases - sql, no-sql, object store</li> <li>● Application and Microservice Security- OAuth, access tokens</li> <li>● Understand Autoscale - horizontal and vertical scaling, logging and monitoring aspects of apps and infrastructure</li> <li>● Learning DevOps frameworks - toolchains, ci/cd, blue/green deployment, canary deployment</li> </ul> <p><b>Module 4 - Case Studies - Public Cloud Provider – aws, azure, ibmcloud</b></p>
4	<b>Texts/References</b>	<p>Text Books:</p> <ul style="list-style-type: none"> <li>- Thomas Erl, Zaigham Mahmood, Ricardo Puttini, “Cloud Computing Concepts, Technology &amp; Architecture”, Pearson, 2013.</li> </ul> <p>Reference Books:</p> <ul style="list-style-type: none"> <li>- Boris Scholl, Trent Swanson, Peter Jausovec, “Cloud Native”, O’Reilly, 2019.</li> </ul> <p>Resources from Internet:</p> <p>Public Cloud Documentations:</p> <p><a href="https://learning.oreilly.com/library/view/cloud-computing-concepts/9780133387568/">https://learning.oreilly.com/library/view/cloud-computing-concepts/9780133387568/</a></p> <p><a href="https://www.amazon.in/Cloud-Computing-Concepts-Technology-Architecture/dp/0133387526/">https://www.amazon.in/Cloud-Computing-Concepts-Technology-Architecture/dp/0133387526/</a></p> <p>Class Notes/Lectures</p>

1	<b>Title of the course</b> (L-T-P-C)	<b>Programming Parallel Machines</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Good programming experience in C, C++ or Fortran.
3	<b>Course content</b>	Elements of parallel algorithm design (domain decomposition, mapping), Parallel hardware architecture - overview of shared-memory systems, cache coherence, and sequential consistency. Shared- memory programming (OpenMP, PThread, Cilk). Scalable algorithm design and measuring performance (speedup, efficiency). Distributed-memory programming (MPI and message-passing model, global address space model). Interconnects. Introduction to GPU programming
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Ananth Grama, Anshul Gupta, George Karypis, Vipin Kumar: Introduction to Parallel Computing, Addison Wesley 2003.</li> <li>2. <a href="https://computing.llnl.gov/tutorials/mpi/">https://computing.llnl.gov/tutorials/mpi/</a></li> <li>3. <a href="https://computing.llnl.gov/tutorials/open_MP/">https://computing.llnl.gov/tutorials/open_MP/</a></li> </ol>

1	<b>Title of the course</b> (L-T-P-C)	<b>Introduction to Abstractions and Paradigms for Programming</b> <b>(1-0-0-2)</b>
2	<b>Pre-requisite courses(s)</b>	
3	<b>Course content</b>	Ideas behind imperative, applicative, object oriented and logic programming paradigms such as typing, expressions, pure functions, recursion, higher order functions, encapsulation, inheritance, goal satisfaction, backtracking, unification.
4	<b>Texts/References</b>	Harold Abelson, Gerald Jay Sussman and July Sussman, Structure and Interpretation of Computer Programs, 2nd edition, The MIT Press, 1996. David A. Watt, Programming Language Concepts and Paradigms, Prentice-Hall, 1990. Rajeev Sangal, Programming Paradigms in Lisp, McGraw Hill, 1991.

1	<b>Title of the course</b> (L-T-P-C)	<b>Operational Analysis</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Exposure to programming languages (C/C++/MATLAB), ability to rig up basic electrical/electronic circuits
3	<b>Course content</b>	<p>Part 1: Stages of manufacturing supply chain and value addition. Concepts of supply, storage, production, warehousing and transport. Supply chain terminology. Structure of supply chains, decisions and decision levels in supply chain. Overview of supply chain operations &amp; costs: location, procurement, production, inventory, transportation, and information technology. Material flow systems: push, pull, flow shop, job shop, cell, etc.</p> <p>Part 2: Quantitative models in supply chain systems including forecasting, production planning, scheduling and inventory models. Basic forecasting models. Product structure representation. Bill of materials. Material accounting logic and MRP systems. Scheduling and sequencing of parts: single/two machine. Notion of quality &amp; quality control. Quality and yield of manufacturing processes.</p> <p>Inventory models: EOQ, periodic review, continuous review. Introduction to procurement and distribution models. Transport &amp; logistics costs. Project management techniques: CPM and PERT.</p>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. Steven Nahmias (2004) Production and Operations Analysis, 5th edition, McGraw-Hill.</li> <li>2. W. Hopp and M. Spearman (2000) Factory Physics, 3rd edition, McGraw-Hill.</li> <li>3. Sunil Chopra and Peter Meindl (2006) Supply Chain Management, 3rd edition, Pearson Education India.</li> <li>4. Joseph S. Martinich (2000) Production and Operations Management: An Applied Modern Approach, John Wiley &amp; Sons.</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Advanced in Cloud Technologies</b>
2	<b>Pre-requisite courses(s)</b>	-
3	<b>Course content</b>	<p><b>Chapter-1: Introduction to cloud (2 Hours)</b></p> <ul style="list-style-type: none"> <li>a. Public</li> <li>b. Private</li> <li>c. On-Prem</li> <li>d. Managed cloud services: IaaS/PaaS/SaaS</li> <li>e. Architectures for the cloud software <ul style="list-style-type: none"> <li>i. Monoliths Vs Microservices</li> <li>ii. Serverless Vs Servers</li> </ul> </li> <li>f. Cloud nativity: <ul style="list-style-type: none"> <li>i. Stateful Vs Stateless</li> <li>ii. Platform independence and scale (designing with chaos in mind)- <a href="https://12factor.net/">https://12factor.net/</a></li> <li>iii. Infrastructure-as-code</li> <li>iv. DevOps-(cloud-native CI/CD)</li> <li>v. Multi-cloud and Hybrid-cloud</li> </ul> </li> </ul> <p><b>Chapter 2: Virtualization(4 Hours)</b></p> <ul style="list-style-type: none"> <li>a. What is Virtualisation? (Bare metal Servers)</li> <li>b. Why virtualisation?</li> <li>c. Virtualization technologies: <ul style="list-style-type: none"> <li>i. Virtual Machines <ul style="list-style-type: none"> <li>Hypervisors</li> <li>Practical Illustration: libvirt API</li> </ul> </li> <li>ii. Container <ul style="list-style-type: none"> <li>Container Image Layering architecture</li> <li>o Copy-on-write(e.g., Union File System)</li> <li>o namespaces vs. cgroups for resource access</li> <li>Hands-On: Docker/Podman Usage tutorials</li> <li>o Dockerfiles</li> <li>o Docker/Podman commands</li> </ul> </li> <li>iii. Emerging technologies <ul style="list-style-type: none"> <li>Unikernels-some demo illustrations</li> <li>Web Assembly- some demo illustrations</li> <li>Secure virtualization <ul style="list-style-type: none"> <li>o Secure containers - gVisor (<a href="https://gvisor.dev/">https://gvisor.dev/</a>)</li> <li>o Kata (<a href="https://ubuntu.com/kubernetes/docs/kata">https://ubuntu.com/kubernetes/docs/kata</a>)</li> <li>o Nabla (<a href="https://nabla-containers.github.io/">https://nabla-containers.github.io/</a>)</li> </ul> </li> <li>MicroVM- Firecracker</li> </ul> </li> </ul> </li> </ul> <p><b>Chapter3: Orchestration Platforms (4 hours)</b></p> <ul style="list-style-type: none"> <li>a. Why we need orchestration?</li> <li>b. Popular or orchestration frameworks <ul style="list-style-type: none"> <li>i. Docker swarm</li> <li>ii. Kubernetes <ul style="list-style-type: none"> <li>1. Declarative approach</li> <li>2. Abstractions and Resources</li> <li>3. Auto-Scaling</li> <li>4. Flavours of K8s <ul style="list-style-type: none"> <li>o K8s for the edge-K3s, KubeEdge, MicroK8s, Single-node OCP</li> <li>o virtual Kubelet</li> </ul> </li> <li>5. Multi-cluster <ul style="list-style-type: none"> <li>o Hypershift</li> <li>o KCP- Kubernetes control plane-creating virtual clusters</li> </ul> </li> <li>7. Hands-On <ul style="list-style-type: none"> <li>o Resource YAMLS – pods, deployments, service, ingress, jobs, statefulset &amp; headless service, volumes, configmaps and secrets etc.</li> </ul> </li> </ul> </li> </ul> </li> </ul>

		<p>O Helmchats</p> <ol style="list-style-type: none"> <li>i. Serverless: scale-to-zero       <ol style="list-style-type: none"> <li>1. Function-as-a-service (FaaS)</li> <li>2. Function-as-a-service (FaaS)</li> </ol> </li> </ol> <p><b>Chapter-4: Modernisation (3 Hours)</b></p> <ol style="list-style-type: none"> <li>a. Assessment and analysis (e.g., Konveyor Move2Kube)</li> <li>b. Right-sizing of target cluster</li> <li>c. Approaches to modernization       <ol style="list-style-type: none"> <li>i. Lift-and-shift           <ol style="list-style-type: none"> <li>1. Kubevirt:               <p><a href="https://www.youtube.com/watch?v=-w4Afj5-0_g">https://www.youtube.com/watch?v=-w4Afj5-0_g</a> [<a href="https://kubevirt.io/">https://kubevirt.io/</a>]</p> </li> <li>ii. Re-platform: Platform modernization                   <ol style="list-style-type: none"> <li>1. What?</li> <li>2. Why?</li> <li>3. Challenges</li> <li>4. Real-world tool: IBM Move2Kube                       <ol style="list-style-type: none"> <li>1. Demo</li> <li>2. Tutorials</li> </ol> </li> </ol> </li> <li>iii. Refactor: Application modernization                   <ol style="list-style-type: none"> <li>1. What?</li> <li>2. Why?</li> <li>3. Challenges</li> <li>4. Real-world tool: IBM Mono2Micro                       <ol style="list-style-type: none"> <li>1. Demo</li> <li>2. Tutorials</li> </ol> </li> </ol> </li> </ol> </li> </ol> <p><b>Chapter 5: Operations (1 Hour)</b></p> <ol style="list-style-type: none"> <li>a. DevSecOps - Compliance, Security and Governance</li> <li>b. Application Resource Management (e.g., IBM Turbonomic) Observability (e.g., IBM Instana)</li> </ol> </li></ol>
4	Texts/References	<ol style="list-style-type: none"> <li>1. Konveyor Suite Of Tools: <a href="https://konveyor.github.io/">https://konveyor.github.io/</a></li> <li>2. Jonas, Eric, Johann Schleier-Smith, Vikram Sreekanti, Chia-Che Tsai, Anurag Khandelwal, Qifan Pu, Vaishaal Shankar et al. "Cloud programming simplified: A Berkeley view on serverless computing." <i>arXiv preprint arXiv:1902.03383</i> (2019).</li> <li>3. "Kubernetes Patterns". <a href="https://www.oreilly.com/library/view/kubernetes-patterns/9781492050278/">https://www.oreilly.com/library/view/kubernetes-patterns/9781492050278/</a></li> <li>4. Seshadri, Padmanabha V., Harikrishnan Balagopal, Akash Nayak, Ashok Pon Kumar, and Pablo Loyola. "Konveyor Move2Kube: A Framework For Automated Application Replatforming." In <i>2022 IEEE 15th International Conference on Cloud Computing (CLOUD)</i>, pp. 115-124. IEEE, 2022.</li> <li>5. Talks on Move2Kube [2020-2022]:<a href="https://www.youtube.com/playlist?list=PL4aUFFbk56EMwb9xoYs-3RjgkD38C6nPO">https://www.youtube.com/playlist?list=PL4aUFFbk56EMwb9xoYs-3RjgkD38C6nPO</a></li> <li>Security as code, Published 2023:<a href="https://learning.oreilly.com/library/view/security-as-code/9781098127459/">https://learning.oreilly.com/library/view/security-as-code/9781098127459/</a></li> <li>6. Cloud computing: Methodology, Systems and Applications, Published 2017 <a href="https://learning.oreilly.com/library/view/cloud-computing/9781439856420/">https://learning.oreilly.com/library/view/cloud-computing/9781439856420/</a></li> <li>7. What is Kubevirt? Orchestrating VMs in Kubernetes, Published 2022:<a href="https://learning.oreilly.com/library/view/what-is-kubevirt/9781098133429/">https://learning.oreilly.com/library/view/what-is-kubevirt/9781098133429/</a></li> <li>8. Infrastructure as code, 2<sup>nd</sup> edition, Published 2020:<a href="https://learning.oreilly.com/library/view/infrastructure-as-code/9781098114664/">https://learning.oreilly.com/library/view/infrastructure-as-code/9781098114664/</a></li> <li>9. "Operating Systems: Three Easy Pieces". <a href="https://pages.cs.wisc.edu/~remzi/OSTEP/">https://pages.cs.wisc.edu/~remzi/OSTEP/</a></li> </ol>

		<p>10. "Professional Linux Kernel Architecture". <a href="https://www.oreilly.com/library/view/professional-linux-kernel/9780470343432/">https://www.oreilly.com/library/view/professional-linux-kernel/9780470343432/</a></p> <p>12. Armbrust, Michael, Armando Fox, Rean Griffith, Anthony D. Joseph, Randy H. Katz, Andrew Konwinski, Gunho Lee et al. <i>Above the clouds: A berkeley view of cloud computing</i>. Vol. 17. Technical Report UCB/EECS-2009-28, EECS Department, University of California, Berkeley, 2009.</p> <p>13. <a href="https://www.udemy.com/course/welcome-to-cloud-computing-world/">https://www.udemy.com/course/welcome-to-cloud-computing-world/</a> (Cloud Computing Fundamentals)</p> <p>14. <a href="https://www.udemy.com/course/cloud-computing-for-beginners-infrastructure-as-a-service/">https://www.udemy.com/course/cloud-computing-for-beginners-infrastructure-as-a-service/</a> (Cloud Computing for Beginners - Infrastructure as a Service)</p> <p>15. <a href="https://www.coursera.org/learn/introduction-to-cloud">https://www.coursera.org/learn/introduction-to-cloud</a> (Introduction to Cloud Computing)</p> <p>16. <a href="https://www.coursera.org/learn/cloud-computing-foundations-duke">https://www.coursera.org/learn/cloud-computing-foundations-duke</a> (Cloud Computing Foundations)</p> <p>17. <a href="https://www.coursera.org/specializations/cloud-computing">https://www.coursera.org/specializations/cloud-computing</a> (Cloud Computing Specialization)</p> <p>18. <a href="https://www.coursera.org/learn/cloud-computing">https://www.coursera.org/learn/cloud-computing</a> (Cloud Computing Concepts, Part 1)</p> <p>19. <a href="https://www.coursera.org/learn/cloud-computing-2">https://www.coursera.org/learn/cloud-computing-2</a> (Cloud Computing Concepts: Part 2)</p> <p>20. <a href="https://www.coursera.org/learn/fundamentals-of-cloud-computing">https://www.coursera.org/learn/fundamentals-of-cloud-computing</a> (Fundamentals of Cloud Computing)</p> <p>21. <a href="https://www.coursera.org/specializations/cloud#courses">https://www.coursera.org/specializations/cloud#courses</a> (System Issues in Cloud Computing Specialization)</p> <p>22. <a href="https://www.coursera.org/learn/codio-cloud-computing-primer-semi-tech-business-software-as-a-service-saas">https://www.coursera.org/learn/codio-cloud-computing-primer-semi-tech-business-software-as-a-service-saas</a></p> <p>23. <a href="https://www.coursera.org/learn/codio-cloud-computing-primer-semi-tech-business-infrastructure-as-a-service-iaas">https://www.coursera.org/learn/codio-cloud-computing-primer-semi-tech-business-infrastructure-as-a-service-iaas</a></p> <p><a href="https://www.coursera.org/learn/codio-cloud-computing-primer-semi-tech-business-platform-as-a-service-paas">https://www.coursera.org/learn/codio-cloud-computing-primer-semi-tech-business-platform-as-a-service-paas</a></p>
--	--	---

1	<b>Title of the course</b> (L-T-P-C)	<b>Advanced Topics in Deep learning</b> <b>(3-0-2-8)</b>
2	<b>Pre-requisite courses(s)</b>	1. Deep Learning (CS 428) or NNDL (EE620) or equivalent and
3	<b>Course content</b>	<p><b>Module 1: Overview of Deep Learning</b> Historical Context, Basics and Milestones, Basic Optimization, Convex vs. Non-Convex Optimization, Convergence Rates and Acceleration effects of algorithms, Learning vs. Pure Optimization</p> <p><b>Module 2: Optimization Methods and Regularization</b> Advanced Optimization Algorithms, Parameter Initialization Strategies, Algorithms with Adaptive Learning Rates, Approximate Second Order Methods, Parameter Norm Penalties, Norm Penalties as Constrained Optimization, Early Stopping, Dropout, Tangent Distance</p> <p><b>Module 3: Learning Paradigms for Deep Learning</b> Generalization Properties, Efficient Learning, Guarantees on Learning, Teacher-student network, Zero-Shot Learning, Few-Shot Learning, Orthogonal Learning.</p> <p><b>Module 4: Transformer Networks</b> Limitation of Recurrent Neural Networks (RNNs), Sequence-to- sequence modeling and attention, Transformers, Large language models (LLM) – Bidirectional Encoder Representations from Transformers (BERT) and Generative Pre-Trained Transformers (GPT), Transformer for vision and speech applications</p> <p><b>Module 5: Advanced Convolutional Networks</b> Advanced Topics on Transfer Learning, Advanced CNN Architectures and Applications including Object Detection and Recognition, Segmentation (Triplet loss, contrastive loss, and ranking loss), CNN for graph data</p> <p><b>Module 6: Recent Trends in Deep learning</b> Deep Generative AI - Generative Adversarial Networks (GANs) Series, Evolutionary Deep Learning, Deep Reinforcement Learning, Security and Privacy Concerns in Deep Learning</p> <p><b>Module 7: Advanced Emerging Models</b> Attention and Diffusion Models, Physics-Informed Neural Networks, Kolmogorov-Arnold Networks (KAN), Conformers</p> <p>The lab component will closely follow the theory course, allowing students to implement advanced algorithms on different topics studied. Practical sessions will include hands-on experiments with advanced deep learning algorithms through Python programming, covering areas like optimization, regularization techniques, Zeroshot / Fewshot Learning, RNNs, transformers, Advanced CNN Architectures and Applications, Generative AI, Physics informed neural networks, and emerging models.</p>

4	<b>Texts/References</b>	<ul style="list-style-type: none"> <li>• Goodfellow, Y. Bengio, and A. Courville, Deep Learning. MIT Press, 2016.</li> <li>• I. Vasilev, Advanced Deep Learning with Python: Design and Implement Advanced Next-Generation AI Solutions Using TensorFlow and PyTorch. Packt Publishing Ltd, 2019.</li> <li>• Y. Bengio, "Learning deep architectures for AI," Foundations and Trends® in Machine Learning, vol. 2, no. 1, pp. 1-127, 2009.</li> <li>• C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, "Understanding deep learning (still) requires rethinking generalization," Communications of the ACM, vol. 64, no. 3, pp. 107-115, 2021.</li> <li>• B. Neyshabur, R. Tomioka, R. Salakhutdinov, and N. Srebro, "Geometry of optimization and implicit regularization in deep learning," arXiv preprint arXiv:1705.03071, 2017.</li> <li>• C. Jin, R. Ge, P. Netrapalli, S. M. Kakade, and M. I. Jordan, "How to escape saddle points efficiently," in International Conference on Machine Learning, 2017, pp. 1724-1732.</li> <li>• S. Cameron, A. Pretorius, and S. Roberts, "Nonparametric Boundary Geometry in Physics Informed Deep Learning," Advances in Neural Information Processing Systems, vol. 36, 2024.</li> <li>• H. Wang, S. Ma, L. Dong, S. Huang, D. Zhang, and F. Wei, "Deepnet: Scaling transformers to 1,000 layers," IEEE Transactions on Pattern Analysis and Machine Intelligence, 2024.</li> <li>• B. Min, H. Ross, E. Sulem, A. P. B. Veyseh, T. H. Nguyen, O. Sainz, and D. Roth, "Recent advances in natural language processing via large pre-trained language models: A survey," ACM Computing Surveys, vol. 56, no. 2, pp. 1-40, 2023.</li> <li>• W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, and J. R. Wen, "A survey of large language models," arXiv preprint arXiv:2303.18223. 2023.</li> <li>• N. Li, L. Ma, G. Yu, B. Xue, M. Zhang, and Y. Jin, "Survey on evolutionary deep learning: Principles, algorithms, applications, and open issues," ACM Computing Surveys, vol. 56, no. 2, pp. 1-34, 2023.</li> </ul>
---	-------------------------	---

1	<b>Title of the course</b> (L-T-P-C)	<b>Adversarial Machine Learning</b> <b>(3-0-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Nil
3	<b>Course content</b>	<p><b>Unit 1: Introduction to Machine Learning</b> Overview of machine learning, pre-processing required for the audio, image and video classification, construction of data loaders for audio, image and video classification, building of deep learning models for audio and video classification, Recent development of deep learning for audio and video classification.</p> <p><b>Unit 2: Basics of White-box and Black-Box Adversarial Attacks</b> Clean sample for deep models, Significance of adversarial sample, crafting of adversarial samples using white box techniques, advantage and disadvantages of white box attack methods, computational (time and memory) analysis of white box attacks. Construction of adversarial samples using black box attacks, advantage and disadvantage of black box attacks, computational analysis(time, memory) of black box attacks.</p> <p><b>Unit 3: Transferability of Adversarial Attacks</b> Transferability of white box attacks, Transferability of black box attacks, recent trends in enhancing the transferability of adversarial attacks across deep models.</p> <p><b>Unit 4: Defense Against Adversarial Attacks</b> Assessment of robustness of deep models against adversarial attacks, adversarial training of a deep model, Input transformation-based defense techniques against adversarial attacks, encryption-based defense techniques, Failure of adversarial training and challenges.</p> <p><b>Unit 5: Advance Attack Techniques:</b> Poisoning attacks, backdoor attack, backdoor attacks on images, backdoor attack on audio, image and videos. Detection of backdoor attacks. Explaining and understanding the impact of adversarial attacks on deep model using explainability analysis. Explainability assisted adversarial attacks. Explainability as a defense against adversarial attacks.</p>
4	<b>Texts/References</b>	<p><b>Textbook(s):</b></p> <ul style="list-style-type: none"> <li>• David J. Miller, Zhen Xiang, Urbana-Champaign, George Kesidis, Adversarial Learning and Secure AI, Cambridge University Press, 2023</li> <li>• Aneesh Sreevallabh Chivukula, Xinghao Yang, Bo Liu, Wei Liu,</li> <li>• Adversarial Machine Learning: Attack Surfaces, Defense Mechanisms, Learning Theories in Artificial Intelligence, Springer, 2023.</li> </ul> <p><b>Paper (s):</b></p> <ul style="list-style-type: none"> <li>• Chakraborty, Anirban and Alam, Manaar and Dey, Vishal and Chattopadhyay, Anupam and Mukhopadhyay, Debdeep, A survey on adversarial attacks and defenses, CAAI Transactions on Intelligence Technology, Willey Online, 2021</li> <li>• Baniecki, Hubert and Biecek, Przemyslaw, Adversarial attacks and defenses in explainable artificial intelligence: A survey, Information Fusion, Elsevier, 2024.</li> </ul>

1	<b>Title of the course (L-T-P-C)</b>	<b>Machine Learning Applications in Wireless Networks (2-1-0-6)</b>
2	<b>Pre-requisite courses(s)</b>	Nil
3	<b>Course content</b>	<p><b>Unit 1 Introduction to Wireless Networks and Fundamentals of Machine Learning for Wireless Networks:</b> Overview of Wireless Networks Cellular, Wi-Fi, and IoT networks; Introduction to Machine Learning: Basic concepts and types of ML algorithms; Challenges and Opportunities in Wireless Networks; Data Preprocessing for Wireless Networks; Supervised, Unsupervised, and Reinforcement Learning in Wireless Networks; Feature Selection and Engineering</p> <p><b>Unit 2 Network Performance Optimization and QoS improvement:</b> Dynamic Resource Allocation in Wireless Networks; Traffic Management and Load Balancing with ML; Predictive Maintenance for Wireless Network Devices; QoS in Wireless Networks; Bandwidth Allocation and Optimization; Prioritizing Traffic with Machine Learning</p> <p><b>Unit 3 Network planning, design, and resource management</b> Capacity Planning with ML; Topology Optimization for Wireless Networks; Case Studies and Practical Applications; ML for Efficient Resource Allocation; Energy Efficiency in Wireless Networks; Optimization of Network Device Placement</p> <p><b>Unit 4 Network traffic prediction and traffic engineering:</b> Traffic Forecasting using ML; Predictive Analytics for Network Traffic; Adaptive routing models; fault prediction and detection for proactive rerouting of traffic; Real-world Applications and Case Studies</p> <p><b>Unit 5 Wireless Network Security with Machine Learning:</b> Intrusion Detection in Wireless Networks; Malware Detection and Prevention; Anomaly Detection using ML; Real-world Applications and Case Studies</p>
4	<b>Texts/References</b>	<p>Machine Learning and Wireless Communications, by Andrea Goldsmith et al., Cambridge University Press, 2022</p> <p>Machine Learning for Future Wireless Communications, by Fa-Long Luo, Wiley-IEEE press, 2020</p>

1	<b>Title of the course</b> (L-T-P-C)	<b>Cryptography and Network Security</b> <b>2-1-0-6</b>
2	<b>Pre-requisite courses(s)</b>	--
3	<b>Course content</b>	<p><b>CS 204: Computer Networks</b> or Equivalent, and  <b>CS 212: Computer Networks Lab</b> or Equivalent, and  <b>CS 203: Discrete Structures</b> or Equivalent  * Equivalent is decided by instructor</p> <ul style="list-style-type: none"> <li>● Introduction: Threats, Vulnerabilities, Attacks, Data Integrity, Confidentiality, Anonymity, Message and Entity Authentication, Authorization, Non-repudiation</li> <li>● Mathematical Background for Cryptography: Modulo Arithmetic, GCD (Euclid's Algorithm), Algebraic Structures (Groups), Chinese Remainder Theorem</li> <li>● Secret Key Cryptography: DES (Data Encryption Standard), MAC (Message Authentication Codes) and other applications, Attacks</li> <li>● Public Key Cryptography: RSA, ECC, Performance, Practical Issues, PKCS (Public Key Cryptography Standard)</li> <li>● Integrity, Authentication and Non-Repudiation: Cryptographic Hash: Properties, Applications, Performance, Birthday Attack, Hash Functions (Examples - MD5, SHA5) : Properties and implementation, Message Authentication Code (MAC), Digital Signature (RSA, DSA Signatures)</li> <li>● Key Exchange: Discrete Logarithm Problem, Diffie- Hellman Key Exchange, Attacks</li> <li>● Public Key Infrastructure: Digital Certificates, Certification Authorities</li> <li>● Protocols: Basic Authentication Protocols (Attacks(Replay, Reflection, Man-in-the-middle), Needham Schroeder Protocol, Kerberos), Security Protocols (Network Security with IPsec, SSL/TLS)</li> <li>● Software Vulnerabilities: Buffer Overflow, Cross Site Scripting, SQL Injection</li> <li>● Attacks: Attacks on DNS, Phishing and Pharming, Denial of Service (DoS and DDoS)</li> <li>● Intrusion detection system (Host based and Network based)</li> <li>● Honeypots</li> </ul>
4	<b>Texts/References</b>	<ol style="list-style-type: none"> <li>1. William Stallings, Cryptography and Network Security: Principles and Practices, 8th Edition, Pearson, June 2022.</li> <li>2. Bruce Schneier, Applied Cryptography: Protocols, Algorithms, and Source Code in C, 20th Edition, Wiley, May 2017.</li> <li>3. Douglas Robert Stinson, Maura Paterson, Cryptography: Theory and Practice, 4th Edition, CRC Press, January 2018.</li> </ol>

1	<b>Title of the course (L-T-P-C)</b>	<b>Scalable Data Mining Lab 0-0-3-3</b>
2	<b>Pre-requisite courses(s)</b>	--
3	<b>Course content</b>	<ul style="list-style-type: none"> <li>• Basic programs of Hadoop MapReduce: Driver code, Mapper code, Reducer code, RecordReader, Combiner, Partitioner</li> <li>• Big data analytics in Spark using PySpark: Installing Apache Spark, Resilient Distributed Dataset (RDD) and basic operations on RDD</li> <li>• Building machine learning model using PySpark and MLlib</li> <li>• Studying performance of gradient descent and normal equations on regression models</li> <li>• Implementing PageRank for webdata</li> <li>• Building recommender systems</li> <li>• Locality sensitive hashing based on various distance measures such as Jaccard Distance, Euclidean Distance</li> <li>• Analysis of social networks graphs with Stanford Network Analysis Platform (SNAP)</li> <li>• Mining Frequent Itemsets for finding association rules</li> <li>• Analysis of data streams</li> </ul>
4	<b>Texts/References</b>	<p>Hadoop: The Definitive Guide, 4th Edition by Tom White Released April 2015 Publisher(s): O'Reilly Media, Inc. ISBN: 9781491901632</p> <p>Spark: The Definitive Guide Big Data Processing Made Simple, 1st. ed. By Bill Chambers and Matei Zaharia Released February 2018. Publisher(s): O'Reilly Media, Inc.</p> <p>Mining of Massive Datasets. 2nd edition. - Jure Leskovec, Anand Rajaraman, Jeff Ullman. Cambridge University Press (2016). <a href="http://www.mmds.org/">http://www.mmds.org/</a> Apache Cassandra's Documentation:</p> <p>Introduction to Data Mining. 2 nd edition – Pang-Ning Tan, Michael Steinbach, Anuj Karpatne and Vipin Kumar, Addison-Wesley Longman Publishing Co., Inc (2018).</p> <p>Adaptive Stream Mining: Pattern Learning and Mining from Evolving Data Streams, 1 st edition. - Albert Bifet, IOS Press (2010), ISBN: 1607500906, 9781607500902</p> <p><a href="https://cassandra.apache.org/doc/latest/">https://cassandra.apache.org/doc/latest/</a> /<a href="https://snap.stanford.edu/snap/">https://snap.stanford.edu/snap/</a></p>

1	<b>Title of the course</b> (L-T-P-C)	<b>Introduction to Open Radio Access Networks</b> <b>2-1-0-6</b>
2	<b>Pre-requisite courses(s)</b>	CS 204: Computer Networks or Equivalent
3	<b>Course content</b>	<p>This course provides a comprehensive introduction to Open Radio Access Network (O-RAN), covering its fundamental principles, architecture, and industry impact. Participants will explore the evolution of mobile communication networks, understand the differences between traditional and open RAN approaches, and gain hands-on experience deploying O-RAN using open-source tools. Evolution of RAN and Introduction to Open RAN:</p> <ul style="list-style-type: none"> <li>• Overview of mobile network generations (1G to 5G and beyond)</li> <li>• Evolution of RAN architectures: <ol style="list-style-type: none"> <li>1. Distributed RAN (D-RAN)</li> <li>2. Cloud RAN (C-RAN)</li> <li>3. Virtual RAN (vRAN)</li> <li>4. Open RAN (O-RAN)</li> </ol> </li> <li>• Traditional RAN vendors and components: <ol style="list-style-type: none"> <li>1. Remote Radio Unit (RRU)</li> <li>2. Baseband Unit (BBU)</li> <li>3. Interfaces and interconnections</li> </ol> </li> <li>• Benefits and risks of transitioning to O-RAN</li> </ul> <p><b>Open RAN Standardization and Ecosystem:</b></p> <ul style="list-style-type: none"> <li>• Introduction to the O-RAN Alliance and its role in standardization</li> <li>• Collaboration with industry organizations: TIP (Telecom Infra Project), 3GPP</li> <li>• How O-RAN complements 3GPP and other industry standards Comparison of 3GPP-defined functions/nodes vs. O-RAN-defined functions/nodes</li> <li>• Key global Open RAN initiatives, deployments, and industry efforts</li> </ul> <p><b>O-RAN Architecture and Specifications:</b></p> <ul style="list-style-type: none"> <li>• Overall O-RAN architecture and functional split</li> <li>• Key interfaces in O-RAN architecture: <ol style="list-style-type: none"> <li>1. Need for Open Interfaces</li> <li>2. Fronthaul, Midhaul, and Backhaul connectivity</li> </ol> </li> <li>• O-RAN functional components: <ol style="list-style-type: none"> <li>1. Near-Real-Time and Non-Real-Time RIC (RAN Intelligent Controller)</li> <li>2. O-CU (Centralized Unit) and O-DU (Distributed Unit)</li> <li>3. O-RU (Radio Unit)</li> </ol> </li> <li>• The role of Artificial Intelligence and Machine Learning (AI/ML) in Open RAN <ul style="list-style-type: none"> <li>• Designing x/rApps by leveraging AI/ML Algorithms</li> </ul> </li> </ul> <p>Module 4: Open RAN Implementation and Deployment</p> <ul style="list-style-type: none"> <li>• Introduction to open-source tools for O-RAN deployment: <ol style="list-style-type: none"> <li>1. OpenAirInterface (OAI)</li> <li>2. srsRAN</li> <li>3. O-RAN Software Community (ORAN-SC)</li> <li>4. FlexRIC, EdgeRIC, and other emerging platforms</li> </ol> </li> <li>• Step-by-step deployment of an Open RAN setup</li> <li>• Configuration and integration of O-RAN components</li> <li>• Demonstration of real-world Open RAN use cases</li> </ul> <p>Module 5: Challenges and Future Directions in Open RAN</p> <ul style="list-style-type: none"> <li>• Security considerations in O-RAN networks</li> <li>• Performance optimization and resource management</li> <li>• Scalability and interoperability challenges</li> <li>• Future trends and innovations in O-RAN</li> </ul>

4	<b>Texts/References</b>	Reference template: The fields should be in the following order – Author(s), Title, Volume, Edition, Publisher, Year of publication. <ol style="list-style-type: none"><li>1. Ian C. Wong, Aditya Chopra, Sridhar Rajagopal, Rittwik Jana. Open RAN: The Definitive Guide, 1st Edition, Wiley-IEEE Press, September 2023.</li><li>2. Jyrki T. J. Penttinen, Michele Zarri, Dongwook Kim, Open Ran Explained: The New Era of Radio Networks, 1st Edition, Wiley, August 2024</li><li>3. Vladimir Yanover, Open RAN: Technology and Ecosystem, 1st Edition, Wiley, Jan 2025</li><li>4. O-RAN Alliance, <a href="https://www.o-ran.org/">https://www.o-ran.org/</a></li></ol>
---	-------------------------	---