COCHIN UNIVERSITY OF SCIENCE AND TECHNOLOGY

DEPARTMENT OF COMPUTER SCIENCE

PROGRAMME STRUCTURE & SYLLABUS
[2019 ADMISSIONS ONWARDS]

- M.TECH COMPUTER SCIENCE AND ENGINEERING (Data Science and Artificial Intelligence) [Part-Time]

[Page left blank]

M.Tech Computer Science and Engineering (Data Science and Artificial Intelligence) [Part-Time]

1. Programme Description

This programme is specifically designed and offered in response to the rapidly developing field of data science and AI that are currently gaining unprecedented traction in the industry as well as in the job market worldwide. As such, there will be rich opportunities for specially skilled graduating students to work across multiple domains of the digital economy and participate in enhancing India's global competitiveness. The programme hence aims to produce the next generation of highly skilled post graduates that are required to continue propel the high-value economy growth of our country.

This programme strikes a balance between Computer Science, Statistics and Mathematical Sciences, which enables it to provide a more comprehensive training in terms of the computing aspects of data science. It is also system-and-product driven in that it covers the essentials in software and product development, and it provides students with a large amount of practical trainings.

The programme is structured to equip students with the following:

1. To impart principled understanding of data Science field
2. To understand methods to extract non-obvious and useful patterns from large data sets
3. To acquire skill sets required for a data scientist
4. To build machine learning models for challenging problems

1.2 Number of Seats: 18

1.3 Admission: To be offered in alternate years starting from 2019

1.4 Eligibility requirement for admission:

Candidates for admission should possess:

a) A First Class B.Tech/ BE/ AMIE Degrees in Computer Science & Engineering/Information Technology/Electronics & Communication / Electrical and Electronics / MCA/ Postgraduate degree in Maths/Statistics/Computer Science from any Universities in Kerala or an Examination of any other University/Institution accepted by this University as equivalent thereto, with a minimum of 60% marks / 6.5 CGPA (in 10 point scale or equivalent)

b) Candidates should have minimum of Two Years experience in IT industry of repute/ University/ Institutions recognized by appropriate statutory bodies, after acquiring the qualifying degree.

c) A valid GATE score in appropriate branch.

d) In the absence of GATE qualified candidates Non GATE candidates can be admitted on the basis of a Departmental Admission Test (DAT) conducted by the Department/School.

1.5 Course Duration: 6 Semesters

1.6 Course Timings:

A minimum of 8 hours per week is offered in any of the following time slots.

Week Days : 6 p.m. - 8 p.m. (Monday to Friday)
Weekends : 9 a.m. - 4 p.m. (Saturday, Sunday)

[Page left blank]

2. Course Structure and Syllabus

M.Tech Computer Science and Engineering (Data Science and Artificial Intelligence) [Part-Time]

Semester - I								
Sl. No.	Course code	Course Title	Core/ Elective	Credits	Lec.	Tut	Lab	Marks
1	19-475-0101	Probability and Statistics for Data Science	C	4	4	2	2	100
2	19-475-0102	Artificial Intelligence	C	4	4	2	2	100
		Total for Semester I		8	8	4	4	200

Semester - II								
Sl. No.	Course code	Course Title	Core/ Elective	Credits	Lec.	Tut	Lab	Marks
1	19-475-0201	Foundations of Data Science	C	4	4	2	2	100
2	19-475-0202	Machine Learning Algorithms	C	4	4	2	2	100
		Total for Semester II		8	8	4	4	200

Semester - III								
1	19-475-0301	Probabilistic Graphical Models	C	4	4	2	2	100
2		Elective I	E	4	4	2	2	100
		Total for Semester III		8	8	4	4	200

Electives

19-475-0302 : Image and Video Processing
19-475-0303 : Complex Networks: Theory and Applications
19-475-0304 : Advanced Optimization Techniques

Semester - IV								
Sl. No.	Course code	Course Title	Core/ Elective	Credits	Lec.	Tut	Lab	Marks
1	19-475-0401	Deep Learning Architectures	C	4	4	2	2	100
2		Elective II	E	4	4	2	2	100
		Total for Semester IV		8	8	4	4	200

Electives

19-475-0402 : Natural Language Processing with Deep Learning
19-475-0403 : Real-time Video Analytics
19-475-0404 : Bioinformatics

Semester -V								
Sl. No.	Course code	Course Title	Core/ Elective	Credits	Lec.	Tut	Lab	Marks
1	19-475-0501	Project & Viva Voce	C	16	0	0	13	100
2		Elective III	E	4	4	2	2	300
		Total for Semester V		20	4	2	15	400

Electives

19-475-0502 : Parallel Computing with GPU
19-475-0503 : Mining of Massive Datasets
19-475-0504 : Reinforcement Learning

Semester - VI								
Sl. No.	Course code	Course Title	Core/ Elective	Credits	Lec.	Tut	Lab	Marks
1	19-475-0601	Project & Viva Voce	C	20	0	0	15	400

Total credits for Degree: **72**

[Page left blank]

19-475-0101 Probability and Statistics for Data Science

Core/Elective: Core **Semester:** 1 **Credits:** 4

Course Description:

This course introduces fundamental concepts in probability and statistics from a data-science perspective. The aim is to become familiarized with probabilistic models and statistical methods that are widely used in data analysis.

Course Objectives:

- To introduce the concepts of probability and statistics to data scientists
- To get a clear understanding of statistical inference procedures in estimation and testing
- To understand the connect between statistical theory and statistical practice

Course Content:

Module I
Probability theory: probability spaces, conditional probability, independence – Random variables: discrete and continuous random variables, functions of random variables, generating random variables – Multivariate random variables: joint distributions, independence, generating multivariate random variables, rejection sampling – Expectation: Mean, variance and covariance, conditional expectation

Module II
Random process: definition, mean and autocovariance functions, iid sequences, Gaussian and Poisson process , random walk – Convergence of random process: types of convergence, law of large numbers, Central limit theorem, monte carlo simulation – Markov chains: recurrence, periodicity, convergence, markov-chain monte carlo- Gibbs sampling, EM algorithm, variational inference

Module III
Descriptive statistics: histogram, sample mean and variance, order statistics, sample covariance, sample covariance matrix – Frequentist statistics: sampling, mean square error, consistency, confidence intervals, parametric and non-parametric model estimation

Module IV
Bayesian statistics: Bayesian parametric models, conjugate prior, bayesian estimators – Hypothesis testing: testing framework, parametric testing, permutation test, multiple testing – Mixture models: Gaussian mixture models, multinomial mixture models

Module V
Linear regression: linear models, least-squares estimation, interval estimation in simple linear regression, overfitting – Multiple linear regression models: Estimation of model parameters, MLE – Non linear regression: Non linear least squares, transformation to linear model – Generalized linear models: logistic regression models, Poisson regression

REFERENCES:

1. Michael Mitzenmacher and Eli Upfal; Probability and Computing, 2ed, Cambridge University Press, 2017

2. Alan Agresti, Christine A. Franklin and Bernhard Klingenberg; Statistics: The Art and Science of Learning from Data, 4ed, Pearson, 2017
3. Sheldon M Ross; A First Course in Probability, 10ed, Pearson, 2018
4. Robert V Hogg, Joseph W McKean and Allen T Cralg; Introduction to Mathematical Statistics, 8ed, Pearson, 2018
5. Douglas C Montgomery, Elizabeth A Peck and G Geoffrey Vining; Introduction to Linear Regression Analysis, 5ed, Wiley-Blackwell, 2012

Online Resources:

Course notes of Carlos Fernandez-Granda, DS-GA 1002: Probability and Statistics for Data Science
https://cims.nyu.edu/~cfgranda/pages/DSGA1002_fall17/index.html

19-475-0102 Artificial Intelligence

Core/Elective: Core **Semester:** 1 **Credits:** 4

Course Description:

Artificial Intelligence (AI) is a field that has a long history but is still constantly and actively growing and changing. In this course basics of modern AI as well as some of the representative applications of AI along with huge possibilities in the field of AI, which continues to expand human capability beyond our imagination are taught.

Course Objectives:

- To introduce the foundational principles of AI that drive real world complex applications and practice implementing some of these systems
- To equip students with the tools to tackle new AI problems that they may encounter in life

Course Content:

Module I
Introduction: Overview and Historical Perspective-Intelligent Agents-Problem Solving by searching-State Space Search: Depth First Search, Breadth First Search, DFID-Informed search & exploration-Heuristic Search-Best First Search-Hill Climbing-Beam Search-Tabu Search-Randomized Search:Simulated Annealing, Genetic Algorithms-Constraint Satisfaction Problems

Module II
Finding Optimal Paths: Branch and Bound, A*, IDA*, Divide and Conquer approaches-Beam Stack Search-Problem Decomposition: Goal Trees, AO*, Rule Based Systems -Game Playing: Minimax Algorithm, Alpha-Beta Algorithm, SSS*

Module III
Knowledge and reasoning: Propositional Logic-First Order Logic-Soundness and Completeness -Forward and Backward chaining-Resolution-semantic networks-Handling uncertain knowledge–Probabilistic Reasoning –making simple and complex decisions

Module IV
Planning : Planning problems -Planning with state space search -Partial order planning -Planning Graphs –Planning with Propositional logic-Hierarchical planning -Multi agent planning

Module V
Learning: Forms of learning-Inductive learning -Learning decision trees -Explanation based learning -Statistical learning -Instantance based learning –Neural networks-Reinforcement learning

REFERENCES:

1. Stuart Russell and Peter Norvig. Artificial Intelligence: A Modern Approach, 3e, Prentice Hall, 2009
2. Deepak Khemani. A First Course in Artificial Intelligence, 1e, McGraw Hill Education, 2017
3. Stefan Edelkamp and Stefan Schroedl. Heuristic Search: Theory and Applications, 1e, Morgan Kaufmann, 2011
4. Zbigniew Michalewicz and David B. Fogel. How to Solve It: Modern Heuristics, Springer; 2e, 2004
5. Elaine Rich and Kevin Knight. Artificial Intelligence, 3e, Tata McGraw Hill, 2017
6. Patrick Henry Winston. Artificial Intelligence, 1e, Pearson, 2002

<u>**19-475-0201 Foundations of Data Science**</u>

Core/Elective: Core **Semester:** 2 **Credits:** 4

Course Description:

While traditional areas of computer science remain highly important, increasingly researchers of the future will be involved with using computers to understand and extract usable information from massive data arising in applications, not just how to make computers useful on specific well-defined problem. This course introduce the statistics and computer science concepts required to master data science as a subject.

Course Objectives:

- To introduce the mathematical foundations to deal with high dimensional data
- To introduce concepts like random graphs, random walks, markov chains
- To understand basic underpinnings of machine learning algorithms

Course Content:

Module I
High dimensional space: Law of large numbers, geometry of high dimensions, properties of the unit ball, Gaussians in high dimension, random projection and Johnson-Lindenstrauss Lemma, seperating Gaussians – Singular Value Decomposition: Power method to compute SVD, singular vectors and Eigen vectors, Applications of SVD

Module II
Random Graphs: G(n,p) model, phase transitions, giant component, branching process, cycles and full connectivity – Growth models of Random Graphs: Growth models with and without preferential attachment, small world graphs

Module III
Random walks and Markov chains: Stationary distribution, MCMC, Gibbs sampling, areas and volumes, convergence of random walks, random walks in Euclidean space, web as a Markov chain

Module IV
Learning and VC dimention: Linear Separators, the Perceptron Algorithm, and Margins, Nonlinear Separators, Support Vector Machines, and Kernels, Strong and Weak Learning – Boosting – Vapnik-Chervonenkis dimention: Examples of Set Systems, The Shatter Function, The VC Theorem, Simple Learning

Module V
Algorithms for Massive Data Problems: Locality-Sensitive Hashing - shingling of documents, min-hashing. Distance measures, nearest neighbors, frequent itemsets- LSH families for distance measures, Applications of LSH- Challenges when sampling from massive data

Frequency Moments of Data Streams, Counting Frequent Elements, Matrix Algorithms Using Sampling, Sketch of a Large Matrix, Sketches of Documents

REFERENCES:

1. Avrim Blum, John Hopcroft, Ravindran Kannan; Foundations of Data Science, 2018
https://www.cs.cornell.edu/jeh/book.pdf

2. Jure Leskovec, Rajaraman, A., & Ullman, J. D., Mining of Massive Datasets, Cambridge University Press, 2e, 2016
3. Charu C. Aggarwal, Data Streams: Models and Algorithms, 1e, Springer, 2007
4. Michael I Jordan et.al , Frontiers in Massive Data analysis, 1e, National Academies Press, 2013
5. Nathan Marz & James Warren, Big Data: Principles and best practices of scalable realtime data systems, Manning Publications, 2015

<u>19-475-0202 Machine Learning Algorithms</u>

Core/Elective: Core **Semester:** 2 **Credits:** 4

Course Description:

Machine learning is programming computers to optimize a performance criterion using example data or past experience. This course is to discuss many methods that have their bases in different fields: statistics, pattern recognition, neural networks, artificial intelligence, signal processing, control, and data mining. Major focus of the course is on the algorithms of machine learning to help students to get a handle on the ideas, and to master the relevant mathematics and statistics as well as the necessary programming and experimentation

Course Objectives:

- To understand basics to advanced concepts of Machine Learning
- To attain certain amount of statistical and mathematical sophistication to deal with the subject
- To gain confidence in building Machine Learning algorithms and applications
- To understand the multi-disciplinary aspect of the subject

Course Content:

MODULE I

Machine Learning – Examples of Machine Learning applications – Supervised Learning: Learning a class from examples – Learning multiple classes – Regression – Model selection – Bayesian Decision Theory: Classification – Discriminant functions – Association rules – Parametric methods: MLE – Baye's estimator – Parametric classification – Tuning model complexity

MODULE II

Multivariate Methods – Classification – Regression – Dimensionality reduction: LDA – PCA – Factor Analysis – ICA – Locally Linear Embedding – MDS- Probabilistic Learning: Gaussian Mixture Models- EM algorithm- Nearest Neighbor Methods – Distance Measures

MODULE III

Support Vector Machines: Optimal separation – Kernels – SVM algorithm – Extensions to SVM – Optimization and Search: Least-squares optimization – conjugate gradients – Search: Search techniques – Exploitation and exploration – Simulated annealing

MODULE IV

Learning with trees: Decision trees – CART – Ensemble Learning: Boosting – Bagging – Random Forests – Unsupervised Learning: K-Means algorithm – Vector quantization – SOM algorithm – Markov Chain Monte Carlo Methods

MODULE V

Graphical Models: Bayesian Networks – Markov Random Fields – HMMS – Tracking Methods – Deep Belief Networks: Hopfield Network – Boltzmann Machine – RBM – Deep Learning

REFERENCES:

1. Ethem Alpaydin, Introduction to Machine Learning, 3e, MIT Press, 2014
2. Tom M. Mitchell, Machine Learning, McGraw Hill Education; 1e, 2017
3. Stephen Marsland, Machine Learning, An Algorithmic Perspective, 2e, CRC Press, 2015
4. Giuseppe Bonaccorso, Machine Learning Algorithms, 1e, Packt Publishing Limited, 2017
5. Ethem Alpaydin, Machine Learning- The New AI, MIT Press, 1e, 2016

ONLINE RESOURCES

1. Rohit Singh, TommiJaakkola, and Ali Mohammad.*6.867 Machine Learning.* Fall 2006. Massachusetts Institute of Technology: MIT OpenCourseWare, https://ocw.mit.edu
2. Andrew Ng, https://www.coursera.org/learn/machine-learning

19-475-0301 Probabilistic Graphical Models

Core/Elective: Core **Semester:** 3 **Credits:** 4

Course Description:

Probabilistic graphical models (PGM) is one of the most advanced techniques in machine learning to represent data and models in the real world with probabilities. PGM present a general framework for constructing and using probabilistic models of complex systems that would enable a computer to use available information for making decisions. This course is for anyone who has to deal with lots of data and draw conclusions from it, especially when the data is noisy or uncertain. Data scientists, machine learning enthusiasts, engineers, and those who curious about the latest advances in machine learning will find PGM interesting

Course Objectives:

- Understand the concepts of PGM and which type of PGM to use for which problem
- To understand techniques for representation, inference and learning from graph based models
- To apply Bayesian networks and Markov networks to many real world problems

Course Content:

Module I
Probabilistic reasoning: Representing uncertainty with probabilities – Random variables and joint distributions – Independence – Querying a distribution - Graphs

Module II
Representation: Bayesian Network (BN) representation – Independencies in BN – Factorizing a distribution – D-separation- Algorithm for D-separation – From distributions to Graphs

Module III
Undirected Graphical Models: Factor products – Gibbs distribution and Markov networks – Markov network independencies – Factor graphs – Learning parameters – Conditional Random Fields

Module IV
Gaussian Network Models: Multivariate Gaussians – Gaussian Bayesian networks – Gaussian Markov Random Fields – Exact Inference: variable elimination- Sum-product and belief updates – The Junction tree algorithm

Module V
Learning: Learning Graphical Models – Learning as optimization – Learning tasks – Parameter estimation – Structure learning in BN – Learning undirected models – Actions and decisions

REFERENCES:

1. Daphne Koller, Nir Friedman, Probabilistic Graphical Models- Principles and Techniques, 1e, MIT Press, 2009
2. Christian Borgelt, Rudolf Kruse and Matthias Steinbrecher, Graphical Models- Methods for data analysis and Mining, 2e, Wiley, 2009
3. David Bellot, Learning Probabilistic Graphical Models in R, Packt Publishing, 1e, 2016
4. Luis Enrique Sucar, Probabilistic Graphical Models, 1e, Springer Nature, 2015

19-475-0302: IMAGE AND VIDEO PROCESSING

Core/Elective: Elective **Semester:** 3 **Credits:** 4

Course Description:

The aim of this course is to inculcate a comprehensive knowledge about various Digital Image and Video Processing techniques.

Course Objectives:

- Give an in-depth knowledge about the basic theory and algorithms related to Digital Image and Video Processing.
- Provide awareness about the current technologies and issues specific to Digital Image and Video Processing.
- Provide hands-on experience in using computers to process digital images and Videos.
- Expose students to Python and OpenCV library to do image and video processing tasks.

Course Content:

Module I
Signals: Impulse Sequence - Exponential Sequence - Periodic Sequence. Linear Systems - Shift-Invariant systems - Linear Shift Invariant (LSI) systems – Convolution - Correlation. Image Transforms: Fourier Transform - Discrete Fourier Transform - Z-transform – KL Transform. Causal Systems - Random Signals - Stationary Process - Markov Process.

Module II
Intensity Transformation and Spatial Filtering: Intensity Transformation Functions. Histogram Processing: Histogram Equalization - Histogram Matching. Image enhancement: Arithmetic/Logic operations - Image Subtraction - Image Averaging. Spatial Filtering: Smoothening Spatial Filters - Sharpening Spatial Filters - Laplacian Filter - Unsharp masking - High Boost Filter. Gradient operators: Edge detection filters. Frequency Domain Smoothening - Frequency Domain Sharpening Filters - Laplacian in Frequency domain - Homomorphic Filtering.

Module III
Image degradation/Restoration process model - Noise probability density functions - Spatial Filtering: Mean Filters - Order-statistics filter - Adaptive Filters - Periodic Noise Reduction –Band-reject filters - Band-pass filters - Notch filters. Inverse filtering - Wiener filtering - Performance measures. Color image processing: Color fundamentals - Color models – RGB, CMYK – HIS - Color image smoothening and sharpening – Color image histogram - Color edge detection.

Module IV
Point and line detection - Hough Transform. Image Segmentation: Fundamentals – Thresholding – Otsu's optimum global thresholding - Region-based segmentation: Region growing - Region Splitting and Merging - Segmentation using Morphological Watersheds.

Module V
Color video processing: Video display - Composite versus component video - Progressive and interlaced scan. Motion estimation: Optical flow - pixel based motion estimation - block matching algorithm - deformable block matching algorithm - Global and region based motion estimation - multi-resolution motion estimation - Feature based motion estimation. Stereo and multi-view sequence processing: Depth perception - Stereo imaging principle - Disparity estimation.

REFERENCES:

1. Rafael C. Gonzalez, Richard E. Woods, "Digital Image Processing", 4th Ed., Pearson, March 2017.
2. Anil K. Jain, “Fundamentals of Digital Image Processing”, Pearson, 1st Ed., 1988.
3. William K. Pratt, “Digital Image Processing: PIKS Scientific Inside”, John Wiley & Sons, 4th Ed., 2007.
4. Azriel Rosenfield, Avinash C. Kak, "Digital Picture Processing", Morgan Kaufmann, 2nd Ed., 1982.
5. Bernd Jahne, “Digital Image Processing”, Springer, 6th Ed., 2005.
6. Yao Wang, Jorn Ostermann, Ya-Qin Zhang, "Video Processing and Communications", Pearson, 1st Ed., 2001.
7. Alan C. Bovik, "The Essential Guide to Video Processing", Academic PRess, 2nd Ed., 2009
8. A. Murat Tekalp, "Digital Video Processing", Prentice Hall, 2nd Ed., 2015.

19-475-0303: COMPLEX NETWORKS: THEORY AND APPLICATIONS

Core/Elective: Elective **Semester:** 3 **Credits:** 4

Course description:

Complex networks provide a powerful abstraction of the structure and dynamics of diverse kinds of interaction viz people or people-to-technology, as it is encountered in today's inter-linked world. This course provides the necessary theory for understanding complex networks and applications built on such backgrounds.

Course Objectives

- Representation and analysis of complex networks

Course Content

Module I
Networks of information – Mathematics of networks – Measures and metrics – Large scale structure of networks – Matrix algorithms and graph partitioning

Module II
Network models – Random graphs – walks on graphs - Community discovery – Models of network formation – Small world model - Evolution in social networks – Assortative mixing-
Real networks - Evolution of random network - Watts-Strogatz model – Clustering coefficient - Power Laws and Scale-Free Networks – Hubs - Barabasi-Albert model – measuring preferential attachment- Degree dynamics – non-linear preferential attachment

Module III
Processes on networks – Percolation and network resilience – Epidemics on networks – Epidemic modelling - Cascading failures – building robustness- Dynamical systems on networks – The Bianconi-Barabási model – fitness measurement – Bose-Einstein condensation

Module IV
Models for social influence analysis – Systems for expert location – Link prediction – privacy analysis – visualization – Data and text mining in social networks - Social tagging

Module V
Social media - Analytics and predictive models – Information flow – Modelling and prediction of flow - Missing data - Social media datasets – patterns of information attention – linear influence model – Rich interactions

REFERENCES:

1. Mark J. Newman, Networks: An introduction, 1e, Oxford University Press, 2010
2. Charu C Aggarwal (ed.), Social Network Data Analytics, 1e, Springer, 2011
3. David Easley and Jon Kleinberg, Networks, Crowds, and Markets: Reasoning about a highly connected World, 1e, Cambridge University Press, 2010
4. Albert-Laszlo Barabasi, Network Science, 1e, Cambridge University Press, 2016

19-475-0304: ADVANCED OPTIMIZATION TECHNIQUES

Core/Elective: Elective **Semester:** 3 **Credits:** 4

Course Description:

This course is about the well-known population-based optimization techniques developed during last three decades. This course emphasizing on the advanced optimization techniques to solve large-scale problems especially with nonlinear objective functions

Course Objectives:

- To study concepts of Population based Optimization techniques
- To understand the mathematical foundations for various advanced optimization techniques
- To apply the algorithms to various inter disciplinary applications

Course Content:

Module I

Introduction to optimization- formulation of optimization problems-Review of classical methods-Linear programming-Nonlinear programming-Constraint optimality criteria-constrained optimization-Population based optimization techniques

Module II

Genetic Algorithm-Introduction-Working principle-Representation-selection-fitness assignment-reproduction-cross over-mutation-constraint handling-advanced genetic algorithms-Applications- Artificial Immune Algorithm-Introduction-Clonal selection algorithm- Negative selection algorithm-Immune network algorithms-Dendritic cell algorithms

Module III

Differential Evolution-Introduction-Working principles-parameter selection-advanced algorithms in Differential evolution-Biogeography-Based Optimization-Introduction-Working Principles- Algorithmic variations

Module IV

Particle Swarm Optimization-Introduction- Working principles- Parameter selection- Neighborhoods and Topologies-Convergence-Artificial Bee Colony Algorithm-Introduction- Working principles- Applications-Cuckoo search based algorithm-Introduction- Working principles- Random walks and the step size-Modified cuckoo search

Module V

Hybrid Algorithms-Concepts- divide and conquer- decrease and conquer-HPABC-HBABC-HDABC-HGABC-Shuffled Frog Leaping Algorithm-- Working principles -Parameters- Grenade Explosion Algorithm-Working principle-Applications

REFERENCES:

1. Dan Simon, Evolutionary Optimization Algorithms, 1e, Wiley, 2013
2. Xin-She Yang, Engineering Optimization: An Introduction with Meta-heuristic Applications, 1e, Wiley, 2010
3.S.S. Rao, Engineering Optimization: Theory and Practice, 4e,New Age International, 2013
4.R. VenkataRao, Teaching Learning Based Optimization Algorithm: And Its Engineering Applications, 1e, Springer, 2016

<u>19-475-0401: DEEP LEARNING ARCHITECTURES</u>

Core/Elective: Core **Semester:** 4 **Credits:** 4

Course Description:

Deep learning is part of a broader family of machine learning methods based on learning data representations, as opposed to task-specific algorithms. This course describes deep learning techniques used by practitioners in industry, including deep feed forward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology. This course is useful to students planning careers in either industry or research, and for software engineers who want to begin using deep learning in their products or platforms

Course Objectives:

- To develop a clear understanding of the motivation for deep learning
- To get a practical understanding of machine learning methods based on learning data
- To design intelligent systems that learn from complex and/or large-scale datasets
- To apply deep learning to practical problems

Course Content:

Module I

Deep Networks: Feed forward networks – Learning XOR- Gradient based Learning – Hidden units – Architecture design- Back propagation – Differentiation algorithms

Module II

Regularization for Deep Learning: Penalties-Constrained optimization-Under constrained problems- Dataset augmentation-Semi Supervised learning- Sparse representation- Adversarial training- Optimization for training deep models: Basic algorithms-Algorithms with adaptive learning rates

Module III

Convolutional Networks: Convolution-Pooling-Variants of pooling- Efficient convolutional algorithms – Recurrent and Recursive Nets: Recurrent Neural Networks-Deep Recurrent Networks-Recursive Neural Networks- Explicit memory

Module IV

Linear Factor Models: Probabilistic PCA- ICA – Slow feature analysis – Sparse coding – Autoencoders: UndercompleteAutoencoders – Regularized Autoencoders- Learning Manifolds-Applications of Autoencoders – Representation learning

Module V

Deep generative models: Boltzmann Machines – RBM- Deep Belief Networks-Deep Boltzmann Machines-Convolutional Boltzmann Machines- Directed generative Nets

REFERENCES:

1. Ian Goodfellow, YoshuaBengo, Aaron Courville, Deep Learning, 1e, MIT Press, 2017
2. Nikhil Buduma and Nicholas Locascio, Fundamentals of Deep Learning: Designing Next-Generation Machine Intelligence Algorithms, 1e, Shroff/O'Reilly, 2017
3. Josh Patterson and Adam Gibson, Deep Learning: A Practitioner's Approach, 1e, Shroff/O'Reilly, 2017

19-475-0402 Natural Language Processing with Deep Learning

Core/Elective: Elective **Semester**: 4 **Credits**: 4

Course Description:

Natural language processing (NLP) is one of the most important technologies of the information age. Applications of NLP are everywhere because people communicate mostly everything in language: web search, advertisement, emails, customer service, language translation, radiology reports, etc. Recently, deep learning approaches have obtained very high performance across many different NLP tasks. In this course students will learn to implement, train, debug, visualize and invent their own neural network models. The course provides a deep excursion into cutting-edge research in deep learning applied to NLP

Course Objectives:

- To understand the neural network approach to learn and process natural language data
- To know advanced concepts in natural language processing
- To learn to implement, train, debug and visualize deep neural network models for language processing

Course Content:

Module I

Word Vectors-Singular Value Decomposition- Skip-gram-Continuous Bag of Words (CBOW)-Negative Sampling- Distributed Representations of Words and Phrases and their Compositionality-Efficient Estimation of Word Representations in Vector Space- Advanced word vector representations- language models-softmax-single layer networks

Module II

Neural Networks and backpropagation for named entity recognition-A Neural Network for Factoid Question Answering over Paragraphs-Grounded Compositional Semantics for Finding and Describing Images with Sentences-Deep Visual-Semantic Alignments for Generating Image Descriptions-Recursive Deep Models for Semantic Compositionality over a Sentiment Treebank

Module III

Introduction to Tensorflow- Large-Scale Machine Learning on Heterogeneous Distributed Systems-Recurrent neural networks for language modeling and Extensions of recurrent neural network language model-Opinion Mining with Deep Recurrent Neural Networks

Module IV

GRUs and LSTMs for machine translation- Recursive neural networks for parsing- Parsing with Compositional Vector Grammars-Subgradient Methods for Structured Prediction-Parsing Natural Scenes and Natural Language with Recursive Neural Networks-Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank-Dynamic Pooling and Unfolding Recursive Autoencoders for Paraphrase Detection-Improved Semantic Representations From Tree-Structured Long Short-Term Memory Networks

Module V

Convolutional neural networks for sentence classification- Sequence to Sequence with Neural Networks-Neural Machine Translation by Jointly Learning to Align and Translate- Dynamic Memory Networks for NLP

REFERENCES:

1. Yoav Goldberg, Neural Network Methods for Natural Language Processing, Morgan & Claypool Publishers, 1ed, 2017
2. Ian Goodfellow, YoshuaBengo, Aaron Courville, Deep Learning, 1e, MIT Press, 2017
3. Nikhil Buduma and Nicholas Locascio, Fundamentals of Deep Learning: Designing Next-Generation Machine Intelligence Algorithms, 1e, Shroff/O'Reilly, 2017
4. Josh Patterson and Adam Gibson, Deep Learning: A Practitioner's Approach, 1e, Shroff/O'Reilly, 2017

19-475-0403 Real-time Video Analytics

Core/Elective: Elective **Semester**: 4 **Credits**: 4

Course Description:

This course is about video analytics enabling automated analysis of detection of interesting spatial and temporal events. Image and video analysis include techniques capable of extracting high-level information from the data. Starting from the foundations of image / video analysis this course covers algorithms applied in systems for video analytics so as to develop interesting applications including surveillance

Course Objectives:

- To gain a working knowledge with imgae and video processing
- To understand the analytics on video
- To apply the knowledge to develop applications that use video analytics

Course Content:

Module I:

Fundamentals: Image feature extraction: Feature point detection, Scale Invariant Feature Transform, Edge Detection, Color features. Multiple View Geometry: Perspective Projection Camera Model, Epipolar Geometry, Probabilistic inference, Pattern recognition and Machine learning: SVM and AdaBoost. Background Modeling and Subtraction: Kernel Density Approximation, Background Modeling and Subtraction Algorithms

Module II

Object Detection and Tracking: Pedestrian detection by boosting local shape features: Tree learning algorithms, Edgelet features. Occluded pedestrian detection by part combination. Pedestrian tracking by Associating Detection Responses. Vehicle Tracking and Recognition: Joint tracking and Recognition framework, Joint appearance-motion generative model, Inference algorithm for joint tracking and recognition

Module III

Human Motion Tracking: Image feature representation, Dimension reduction and Movement dynamics learning. Human action recognition: Discriminative Gaussian Process dynamic model. Human Interaction recognition: Learning human activity, Track-body Synergy framework. Multi-camera calibration and global trajectory fusion: Non-overlapping and overlapping cameras. Applications: Attribute-based people search, Soft biometrics for video surveillance: Age estimation from face, Gender recognition from face and body

Module IV

Face Recognition and Gait Analysis: Overview of Recognition algorithms – Human Recognition using Face, Face Recognition from still images, Face Recognition from video, Evaluation of Face Recognition Technologies- Human Recognition using Gait- HMM Framework for Gait Recognition, View Invariant Gait Recognition, Role of Shape and Dynamics in Gait Recognition, Factorial HMM and Parallel HMM for Gait Recognition, Face Recognition Performance

Module V

Behavioral Analysis and Activity Recognition: Event Modeling- Behavioral Analysis- Human Activity Recognition-Complex Activity Recognition- Activity modeling using 3D shape, Video summarization, Shape based activity models, Suspicious Activity Detection. Video Segmentation and Key Frame Extraction: Introduction, Applications of Video Segmentation, Shot Boundary Detection, Pixel-based Approaches, Block-based Approaches, Histogram-based Approaches, Clustering-based Approaches, Performance Measures, Shot Boundary Detection, Key-frame Extraction

REFERENCES:

1. Francesco Camastra, Alessandro Vinciarelli, "Machine Learning for Audio, Image and Video Analysis",Springer Nature, Second Edition, 2015.
2. Yunqian Ma, Gang Qian, “Intelligent Video Surveillance: Systems and Technology”, CRC Press, First Edition, 2009.
3. Fredrik Nilsson, Communications Axis, “Intelligent Network Video: Understanding Modern Video Surveillance Systems”, CRC Press, Second Edition, 2017.
4. Anthony C. Caputo, “Digital Video Surveillance and Security”, Butterworth-Heinemann, Second Edition, 2014.
5. Herman Kruegle, “CCTV Surveillance: Video Practices and Technology”,Butterworth-Heinemann, Second Edition, 2006.
6. Amit K.Roy-Chowdhury, Rama Chellappa, S. Kevin Zhou, Al Bovik, “Recognition of Humans and Their Activities Using Video (Synthesis Lectures on Image, Video, and Multimedia Processing)”, Taxmann Publications Private Limited, 2005.
7. Richard Szeliski, "Computer Vision: Algorithms and Applications", Springer, First Edition, 2010.
8. David A. Forsyth, Jean Ponce, "Computer Vision- A Modern Approach", Pearson Education, Second Edition, 2015.

19-475-0404 Bioinformatics

Core/Elective: Elective **Semester**: 4 **Credits**: 4

Course Description:

This course introduces the study of computational problems in molecular biology. It introduces molecular biology concepts and methodologies for analyzing these data. The course is intended to cover the main aspects which are useful in studying, describing and modeling of Big data problems in the context of molecular biology.

Course Objectives:

- To understand computational problems in molecular biology
- To study molecular structure prediction
- To understand the models of DNA mapping and pathways

Course Content:

Module I

Molecular Biology and Bioinformatics: Introduction to molecular biology- Nucleic acids-DNA-RNA-Proteins-Gene-Genome-Genetic synthesis-Translation-Transcription-Protein synthesis-Chromosomes-Maps and sequences-Human genome project.

Module II

Sequence alignment and database search: Pair-wise sequence alignment- Substitution matrixes -PAM and BLOSSUM matrices, Dot plots - Local and global alignment theory - Dynamic programming methods - FASTA and BLAST algorithms - database search using BLAST and FASTA - Similarity & distance - Similarity scores - Weight matrices - Heuristic method - Hidden Markov Models and their application in sequence analysis

Module III

Phylogenetic trees: Introduction -Dendrogram construction / Molecular Phylogenetics / Tree definitions / Optimality criteria / Distance matrix methods and maximum parsimony / Multiple sequence alignments- tree alignments, star alignments, pattern in pair wise alignment / Genetic algorithm

Module IV

DNA Micro-arrays and Gene Expression- Gene profiling- DNA Microarray technology- Gene regulatory network- Heuristic Algorithms for GRN- S-system model – Computational methods for pathways and system biology- metabolic pathways- genetic pathways- signaling pathways

Module V

Molecular Structure Prediction- RNA secondary structure prediction-Protein Folding problems-Protein threading- Protein structure analysis.

REFERENCES:

1. Zhumar Ghosh and Bibekanand Mallick. Bioinformatics: Principles and Applications., Oxford University Press; 1 edition , March 1, 2015
2. An introduction to Bioinformatics Algorithms 4th Ed: Neil James and Pavel A Pevzner, OUPress, 2014
3. Bioinformatics : Principles and Applications: Zhumur Ghosh, Bibekanand Mallick: OUPress. 2015
4. Building Bioinformatics Solutions: Concord Bessant, Darren Oakley, Ian Shadforth : OU press, 2014
5. Rastogi, S. C., Parag Rastogi, and Namita Mendiratta. Bioinformatics Methods and Applications: Genomics Proteomics and Drug Discovery 4th Ed. PHI Learning Pvt. Ltd., 2013.

<u>19-475-0501 Project & Viva Voce</u>

Core/Elective: Core **Semester**: 5 **Credits**: 16

<u>**19-475-0502 Parallel Computing for Data Science**</u>

Core/Elective: Elective **Semester**: 5 **Credits**: 4

Course Description:

This course is to discuss exclusively on parallel data structures, algorithms, software tools, and applications in data science. It includes examples not only from the classic "n observations, p variables" matrix format but also from time series, network graph models, and numerous other structures common in data science. With the main focus on GPU based computation, the examples illustrate the range of issues encountered in parallel programming.

Course Objectives:

- To gain a working knowledge of parallel programming with data sets
- To develop programming skills required for parallel computing
- To know advanced datastructures required for efficient data processing

Course Content:

Module I

Parallel computing: languages and models for parallelism - Sequential vs parallel: concurrent, parallel, distributed - parallel hardware architecture - modifications to the von Neumann Model - Evolution of GPU - GPGPU - introduction to data parallelism - CUDA program structure - vector addition kernel - device global memory and data transfer

Module II

CUDA thread organization - mapping threads to multi-dimensional data - assigning resources to blocks - synchronization and transparent scalability - thread scheduling and latency tolerance - Memory access efficiency - CUDA device memory types - performance considerations - global memory bandwidth - instruction mix and thread granularity -floating point considerations

Module III

Parallel programming patterns: convolution - prefix sum - sparse matrix and vector multiplication - application case studies - strategies for solving problems using parallel programming

Module IV

Parallel Patterns: merge sort, sequential and parallel approaches, co-rank function implementation, basic parallel merge kernel – Graph search: sequential BFS, parallel BFS, optimizations

Module V

CUDA dynamic parallelism: example for dynamic parallelism, memory data visibility, configurations and memory management, synchronization, streams and events

REFERENCES:

1. David B. Kirk, Wen-mei W Hwu; Programming Massively Parallel Processors, 3 ed, Morgan Kaufmann, 2016
2. Peter Pacheco, An Introduction to Parallel Programming, Morgan Kaufmann, 2011
3. Norman Matloff; Parallel Computing for Data Science, 1 ed, CRC Press, 2015

<u>19-475-0503 Mining of Massive Data Sets</u>

Core/Elective: Elective **Semester**: 5 **Credits**: 4

Course Description:

Big Data concerns large-volume, complex, growing data sets with multiple, autonomous sources. With the fast development of networking, data storage, and the data collection capacity, Big Data is now rapidly expanding in all science and engineering domains. The traditional data mining algorithms also need to be adapted for dealing with the ever-expanding datasets of tremendous volume.

Course Objectives:

- To understand emphasis on the algorithms to be applied on large amounts of data
- To develop hands-on experience on the distributed file systems and MapReduce as a tool for creating parallel algorithms
- To explore streaming data and some of the techniques and algorithms specifically extended for mining on stream data

Course Content:

Module I

Introduction to MapReduce – the map and reduce tasks, MapReduce workflow, fault tolerance. - Algorithms for MapReduce – matrix multiplication, relational algebra operations- Complexity theory for MapReduce

Module II

Locality-Sensitive Hashing - shingling of documents, min-hashing. Distance measures, nearest neighbors, frequent itemsets- LSH families for distance measures, Applications of LSH- Challenges when sampling from massive data

Module III

Mining data streams – stream model, stream data sampling, filtering streams – bloom filters, counting distinct elements in a stream - Flajolet-Martin algorithm. Moment estimates - Alon-Matias-Szegedy algorithm, counting problems for streams, decaying windows

Module IV

MapReduce and link analysis- PageRank iteration using MapReduce, topic-sensitive PageRank - On-line algorithms – Greedy algorithms, matching problem, the adwords problem – the balance algorithm

Module V

Computational model for data mining – storage, cost model, and main memory bottleneck. Hash based algorithm for mining association rule – improvements to a-priori, park-chen-yu algorithm, multistage algorithm, approximate algorithm, limited-pass algorithms – simple randomized algorithm, Savasere, Omiecinski, and Navathe algorithm, Toivonen algorithm

REFERENCES:

1. Jure Leskovec, Rajaraman, A., & Ullman, J. D.; Mining of Massive Datasets, Cambridge India, 2 ed, 2016
2. Charu C. Aggarwal; Data Streams: Models and Algorithms, 1ed, Springer, 2007
3. Michael I Jordan et.al , Frontiers in Massive Data analysis, 1ed, National Academies Press, 2013
4. Nathan Marz & James Warren, Big Data: Principles and best practices of scalable realtime data systems, Manning Publications, 2015

19-475-0504 Reinforcement Learning

Core/Elective: Elective **Semester**: 5 **Credits**: 4

Course Description:

This course aims to introduce the concepts reinforcement learning and to impart an understanding of how reinforcement learning -- along with supervised and unsupervised learning -- form a building block of modern artificial intelligence.

Course Objectives:

- To have a solid understanding of reinforcement learning concepts and where they fit in the machine learning landscape.
- To develop the ability to take a machine learning problem and figure out when it is appropriate to model the problem as a reinforcement learning problem, and how to do that

Course Content:

Module I

Introduction to reinforcement learning: Examples of reinforcement learning, Elements of reinforcement learning - Tabular and Approximate solution methods: Multi-armed bandits, Action-value methods, Incremental Implementation, Upper-Confidence-Bound Action selection, Gradient Bandit Algorithms - Associative Search

Module II

Finite Markov Decision Processes -- The Agent-Environment Interface -- Goals and Rewards -- Returns and Episodes -- Policies and Value Functions -- Optimality of policies and value functions -- Optimality and approximation – Dynamic Programming -- Policy Evaluation -- Policy Improvement -- Policy Iteration -- Value Iteration -- Asynchronous Dynamic Programming -- Generalized Policy Iteration

Module III

Monte Carlo Methods -- Monte Carlo prediction -- Estimation of Action Values -- Monte Carlo Control -- Control without Exploring Starts -- Off-policy prediction via Importance Sampling -- Incremental Implementation -- Off-policy Monte Carlo Control -- Temporal-Difference Learning -- TD prediction -- Advantages of TD methods -- Optimality of TD(0) -- Sarsa and Q-learning -- Expected Sarsa -- Maximization Bias, Double Learning -- Special Cases

Module IV

n-step Bootstrapping -- n-step TD prediction -- n-step Sarsa -- n-step Off-policy Learning -- Off-policy Learning Without Importance Sampling -- Planning and Learning with Tabular Methods -- Models and Planning -- Dyna -- Prioritized Sweeping -- Expected vs.\ Sample Updates -- Trajectory Sampling -- Real-time Dynamic Programming -- Heuristic Search -- Rollout algorithms -- Monte Carlo Tree Search

Module V

Approximate Solution Methods -- On-policy Prediction with Approximation -- Value-function Approximation -- Prediction Objective -- Stochastic-gradient and Semi-gradient Methods -- Linear Methods -- Feature Construction for Linear Methods -- Manual Selection of Step-Size Parameters -- Nonlinear Function Approximation using ANN -- Least-Squares TD -- Memory-based Function Approximation -- Kernel-based Function Approximation -- On-policy Control with Approximation -- Episodic Semi-gradient Control -- Semi-gradient n-step Sarsa -- Average Reward -- Deprecating the Discounted Setting -- Differential Semi-gradient n-step Sarsa

REFERENCES:

1. Richard S.Sutton and Andrew G. Barto; Reinforcement Learning: An Introduction, 2 ed, MIT Press, 2018.
2. Marco Wiering and Martijn van Otterlo (Editors); Reinforcement Learning: State-of-the-Art, Springer, 2012
3. Csaba Szepesvari; Algorithms for Reinforcement Learning, Morgan and Claypool Publishers, 2010

<u>19-475-0601 Project & Viva Voce</u>

Core/Elective: Core **Semester**: 6 **Credits**: 20