Curriculum: MTech in Artificial Intelligence

Centre: Centre of Excellence in Artificial Intelligence, Visleshan I-Hub Foundation

Introduction:

Research in Engineering and Technology have nowadays become multidisciplinary with the phenomenal growth of Artificial Intelligence, Machine Learning, Big Data, Cloud Computing, Robotics, Speech Technologies, Computer Vision, Natural Language Processing etc. Artificial Intelligence is nowadays one of the most discussed technology topics all over in the world. IIT Patna, with generous grant from the Department of Science and Technology (DST) has set up a multidisciplinary centre named Vishlesan I-Hub Foundation under the Technology Innovation Hub (TIH) that targets to leverage Research and Engineering capabilities of Sustainable Development Goals and achieves the mandate of National Mission on Interdisciplinary Cyber Physical Systems, especially in the areas of Video, Speech and Text Analytics. In DPR of TIH, it has been mentioned to start a new M.Tech program, which will provide a platform to create the skilled manpower in the broad areas of Text, Speech and Video Analytics. The earnings in the form of tuition fees collected from this program is conceived to contribute towards making this hub self-sustainable. The TIH under the *Centre of Excellence in Artificial Intelligence* proposes to start a new programme, called M.Tech in “Artificial Intelligence”.

Objectives:

This M.Tech in Artificial Intelligence program will offer students with deep knowledge of core and applied Artificial Intelligence, especially Speech, Video and Text Analytics. This programme is aiming at imparting the necessary breadth and depth to the students for pursuing careers in academics as well as in industry. This programme is aiming at extending undergraduate computing skills with up-to-date and in-depth expertise in specialized areas of Speech Technologies, Computer Vision and Natural Language Processing.

Expected Graduate Attributes:

Students, at the end of this programme, will be able to develop an ability to:

1. Understand the fundamental concepts of Artificial Intelligence, Machine Learning, Big Data, Robotics, Cloud Computing, Speech Processing, Computer Vision, and Natural Language Processing
2. Conceive, Design and Develop state-of-the-art systems, to meet the broad objectives of cyber physical systems
3. Acquire the skills to solve important and practical problems related to speech processing, video analytics and text analytics

Learning Outcomes:

1. Understand the fundamentals concepts of Speech Processing, Video Analytics and Text Analytics

2. Apply appropriate design principles, framework and protocols to develop cyber physical systems.
3. Demonstrate hands-on knowledge of cutting edge speech, video and text analytics tools.
4. Ability to design and develop systems for speech, video and text analytics

Duration: 2 years

Course Fees: Rs. 75,000/Semester
Total intake: 50 (TA-10; RA-10; Self-sponsored: 15; Sponsored-10, Part-time: 5)

Categories: Self-Sponsored, Sponsored, Project-funded (*Research Assistantship*), Regular and Full-time (Teaching Assistantship with *fellowship from the Centre*), Part-time

Eligibility:

- B.Tech./B.E. degree in Computer Science/IT/ECE/Aerospace Engineering/Maths&Computing/ME/Civil Engineering or Equivalent and a valid GATE score in CS/IT, EC, ME, CE
- MSc in Mathematics/Statistics/Mathematics & Computing with valid Gate score in MA or CS/IT, EC, ME, CE.
- **For Sponsored:** 1-2 years of experience

1st SEMESTER

Sl.No.	Course Number	Course Title	L	T	P	C
1	CS561/CS71	Artificial Intelligence / AI Lab-I	3	0	3	9
2	MA501	Probability Statistics and Stochastic Processes	3	0	0	6
3	CS564	Foundations of Machine Learning	3	0	0	6
4	CSXXX	Elective-I	3	0	0	6
5	EE/MEXXX	Elective-II	3	0	0	6
6	HS5XX	HSS Elective	2	0	0	4
TOTAL			**17**	**0**	**3**	**37**

2nd SEMESTER

Sl. No.	Course Number	Course Title	L	T	P	C
1	MA564	Linear Algebra and Optimization techniques	3	0	0	6
2	CS546	Big Data Analytics	3	0	0	6
3	CS551	Intro to Deep Learning	3	0	0	6
4	CS563/EE XXX	Natural Language Processing/ Computer Vision/ Image Processing	3	0	0	6
5	CS514/ME/ EEXXX	Design and Analysis of Algorithms	3	0	0	6
6	CS516	AI Lab-II	0	0	3	3
7	CS589	Comprehensive Viva(After End Semester Exam)	0	0	4	4
TOTAL			**15**	**0**	**4**	**37**

3rd SEMESTER

Sl. No.	Course Number	Course Title	L	T	P	C
1	CS695	Project Thesis-I	0	0	20	20
2.	CS592	Research Seminar	0	0	4	4
TOTAL						**24**

4th SEMESTER

Sl. No.	Course Number	Course Title	L	T	P	C
1	CS696	Project Thesis-II	0	0	24	24
TOTAL						**24**

Total Credit	**37**	**34**	**24**	**24**	**122**

Semester-1: Core Theory

Course No.:MA501	Name: Probability, Statistics and Stochastic Processes	Credits: 3-0-0-6	Prerequisites: NIL

Syllabus:

Algebra of sets, probability spaces, random variables, cumulative distribution functions, mathematical expectations, conditional probability and expectation, moments and inequalities, special discrete and continuous probability distributions, function of a random variable, random vectors and their distributions, convolutions, joint, marginal and conditional distributions, product moments, independence of random variables, bivariate distributions and properties, order statistics and their distributions, sampling distributions, Central Limit Theorem, strong law of large numbers, sequence of random variables, modes of convergence, distributions of the sample mean and the sample variance for a normal population, chi-square, t and F distributions, method of moments and maximum likelihood estimation, concepts of unbiasedness, criteria for choosing estimators, consistency and efficiency of estimates, confidence intervals, pivotal quantities, confidence intervals for proportions, simple and composite hypothesis, null and alternative hypotheses, types of error, level and size of tests, the most powerful test and Neyman - Pearson Fundamental Lemma, tests for one- and two-sample problems for normal populations, tests for proportions, likelihood ratio tests, chi-square test for goodness of fit. discrete and continuous stochastic processes, markov chains, transition probability matrix, state spaces, classification of states, stationary distributions, ergodicity, poisson process, birth and death process.

Books:

References:

1. Rohatgi, V.K., and Saleh, A.K.Md. Ehsanes (2009). An introduction to probability and statistics. Second Edition, Wiley India.
2. Introduction to the Theory of Statistics; Alexander M. Mood, Franklin A. Graybill, Duane C. Boes, Tata McGraw Hill.
3. Milton, J.S. and Arnold, J.C. (2009) Introduction to Probability and Statistics, Fourth Edition, Tata Mcgraw-Hill.
4. Ross, S.M.(2008) Introduction to Probability Models, Ninth edition, Academies Press.
5. Statistical Inference (2007), G. Casella and R.L. Berger, Duxbury Advanced Series .

Course No.:CS561	Name: Artificial Intelligence	Credits: 3-0-0-6	Prerequisites: Nil

Syllabus:

Introduction, Motivation of the course
Problem Solving: Uninformed search, Informed search, local Search, Online search;
Knowledge and Reasoning: Propositional and Predicate Calculus, Semantic Nets, Frames, Scripts, Probabilistic Reasoning
Learning: Introduction to machine learning paradigms: unsupervised, supervised, reinforcement learning, Naive Bayes, Decision Tree, Fundamental of Neural Networks and Deep Learning
Evolutionary Computation: Genetic algorithms, Multi objective optimization, Differential Evolution, Particle Swarm and Ant Colony Optimization
Application Topics: Introduction to NLP, Introduction to Fuzzy Sets and Logic, AI in Social Networks

Books:
References:

1. S. Russel and P. Norvig. Artificial Intelligence: A Modern Approach (Third Edition), Prentice Hall, 2009
2. E. Rich and K. Knight, Artificial Intelligence, Addison Wesley, 1990
3. George Klir, U. St. Clair and B. Yuan, Fuzzy Set Theory: Foundations and Applications, Prentice Hall, 1997
4. Ian Goodfellow, YoshuaBengio and Aaron Courville, Deep Learning, MIT Press, 2016
5. Daphne Koller and Nir Friedman, Probabilistic Graphical Models: Principles and Techniques, MIT Press, 2009.

Journals and conference proceedings

Artificial Intelligence, Machine Learning, ACL Anthology, ICML Proceedings, Proceedings of Uncertainty in AI, ICCV Proceedings and so on.

Course No.:CS564	Foundations of Machine Learning	Credits: 3-0-0-6	Prerequisites: NIL

Syllabus:

Introduction, Logistic regression, Perceptron, Generative learning algorithm: Support vector machines, Model selection and feature selection, Ensemble methods: Bagging, boosting, Random Forest; Unsupervised learning: Clustering: K-means, EM; Mixture of Gaussians, Factor analysis, PCA (Principal components analysis).; Active learning: Theoretical analysis, Committee-based active learning, Active learning from the crowd; Collaborative filtering: Latent factor-based models and neighborhood models; Introduction to Graphical Models (HMM, MEMM, CRF), Deep Learning: CNN, RNN, LSTM, GRU

Texts Books:

1. 1. T. Mitchell. Machine Learning. McGraw-Hill, 1997.
2. 2. Christopher Bishop. Pattern recognition and machine learning. Springer Verlag, 2006.
3. 3. Hastie, Tibshirani, Friedman. The elements of Statistical Learning Springer Verlag.
4. 4. Probability, Random Variables and Stochastic processes by Papoulis and Pillai, 4th Edition, Tata McGraw Hill Edition.
5. 5. A. K. Jain and R. C. Dubes. Algorithms for Clustering Data. Prentice Hall, 198815.

Course No.:CS61/CS571	AI Lab-I/II	Credits: 0-0-3-3	Prerequisites: NIL

Syllabus:

Prolog; Assignment on Logistic regression; Assignment on k-means clustering.

Introduction to Tensorflow, Pytorch, Keras.

Usage of Tensorflow, Pytorch and/or Keras: Simple ML examples; Assignments on NNs; Assignments on CNNs; Assignments on RNN; Assignment on LSTM, GRU

Books:

References

1. Pytorch: https://pytorch.org/assets/deep-learning/Deep-Learning-with-PyTorch.pdf
2. First Contact With TensorFlow: Get Started With Deep Learning Programming By Jordi Torres
3. https://analyticsindiamag.com/top-10-free-books-and-resources-for-learning-tensorflow/
4. https://keras.io/getting_started/learning_resources/
5. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow (second edition), by AurélienGéron

Semester II

Course No.:CS546	Name: Big Data Analytics	Credits: 3-0-0-6	Prerequisites: Nil

Big Data Analytics

Part 1: Introduction to Big Data:

Why Big Data and Where did it come from?, Characteristics of Big Data- Volume, Variety, Velocity, Veracity, Valence, Value, Challenges and applications of Big Data

Part 2: Introduction to Enabling Technologies for Big Data:

Introduction to Big Data Stack, Introduction to some Big Data distribution packages

Part 3: Introduction to Big Data Platforms:

Overview of Apache Spark, HDFS, YARN, Introduction to MapReduce, MapReduce Programming Model with Spark, MapReduce Example: Word Count, Page Rank etc.

Part 4: Introduction to Big Data Storage Platforms for Large Scale Data Storage:

CAP Theorem, Eventual Consistency, Consistency Trade-Offs, ACID and BASE, Introduction to Zookeeper and Paxos, Introduction to Cassandra, Cassandra Internals, Introduction to HBase, HBase Internals

Part 5: Introduction to Big Data Streaming Platforms for Fast Data:

Introduction to Big Data Streaming Systems, Big Data Pipelines for Real-Time computing, Introduction to Spark Streaming, Kafka, Streaming Ecosystem

Part 6: Introduction to Big Data Applications (Machine Learning):

Overview of Big Data Machine Learning, Mahout Introduction, Big Data Machine learning Algorithms in Mahout- kmeans, Naïve Bayes etc.

Part 7: Introduction of Big data Machine learning with Spark:

Big Data Machine Learning Algorithms in Spark- Introduction to Spark MLlib, Introduction to Deep Learning for Big Data

Part 8: Introduction to Big Data Applications (Graph Processing):

Introduction to Pregel, Introduction to Giraph, Introduction to Spark GraphX

Text Books:

Bart Baesens, Analytics in a Big Data World: The Essential Guide to Data Science and its Applications, Wiley, 2014

Reference Book:

1. Dirk Deroos et al., Hadoop for Dummies, Dreamtech Press, 2014.
2. Chuck Lam, Hadoop in Action, December, 2010 | 336 pages ISBN: 9781935182191
3. Mining of Massive Datasets. Leskovec, Rajaraman, Ullman, Cambridge University Press

4. Data Mining: Practical Machine learning tools and techniques, by I.H. Witten and E. Frank

5. Erik Brynjolfsson et al., The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies, W. W. Norton & Company, 2014

Course No.:CS514	**Name: Design and Analysis of Algorithms**	**Credits: 3-0-0-6**	**Prerequisites: NIL**

Syllabus:

Data structures: linked list, stack, queue, tree, balanced tree, graph; Complexity analysis: Big O, omega, theta notation, solving recurrence relation, master theorem
Sorting and searching: Quick sort, merge sort, heap sort; Sorting in linear time; Ordered statistics;
Problem solving strategies: recursion, dynamic programming, branch and bound, backtracking, greedy, divide conquer,
Graph algorithms: BFS, DFS, Shortest path, MST, Network flow;
NP-completeness
Advanced topics: string matching, FFT-DFT, basics of approximation and randomized algorithms;

References:

1. Mark Allen Weiss, "Data Structures and Algorithms in C++", Addison Wesley, 2003.
2. Adam Drozdek, "Data Structures and Algorithms in C++", Brooks and Cole, 2001.
3. Aho, Hopcroft and Ullmann, "Data structures and Algorithm", Addison Welsey, 1984.
4. Introduction to Algorithms Book by Charles E. Leiserson, Clifford Stein, Ronald Rivest, and Thomas H. Cormen

Course No.: CSXXX	**Name:** Deep Learning for Natural Language Processing	**Credits: 3-0-0-6**	**Prerequisites: NIL**

Natural language processing (NLP) is one of the most important technologies of the information age. Understanding complex language utterances is also a crucial part of artificial intelligence. Applications of NLP are everywhere because people communicate most everything in language: web search, advertisement, emails, customer service, language translation, radiology reports, etc. There are a large variety of underlying tasks and machine learning models powering NLP applications. Recently, deep learning approaches have obtained very high performance across many different NLP tasks. These models can often be trained with a single end-to-end model and do not require traditional, task-specific feature engineering. In this spring quarter 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. The final project will involve training a complex recurrent neural network and applying it to a large scale NLP problem. On the model side we will cover word vector representations, window-based neural networks, recurrent neural networks, long-short-term-memory models, recursive neural networks, convolutional neural networks as well as some very novel models involving a memory component. Through lectures and programming assignments students will learn the necessary engineering tricks for making neural networks work on practical problems.

Course Contents:

Intro to NLP and Deep Learning: Linear Algebra, Probability, Optimization and Vector space models

Simple Word Vector representations: word2vec, GloVe: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:GloVe: Global Vectors for Word Representation

Neural Networks and backpropagation: PoS tagging and named entity recognition

Recurrent neural networks -- for language modeling and other tasks: Recurrent neural network based language model, Extensions of recurrent neural network language model, Opinion Mining with Deep Recurrent Neural Networks

Recursive neural networks -- for parsing, Convolutional neural networks -- for sentence classification

Machine Translation, Seq2Seq and Attention

Deep Learning for NLP: Dynamic Memory Networks

Question Answering, Natural Language Generation and Summarization

Contextual Word Representations: BERT

Text and References:

Dan Jurafsky and James H. Martin. Speech and Language Processing (3rd ed. draft)
Jacob Eisenstein. Natural Language Processing
Yoav Goldberg. A Primer on Neural Network Models for Natural Language Processing
Ian Goodfellow, YoshuaBengio, and Aaron Courville. Deep Learning
Delip Rao and Brian McMahan. Natural Language Processing with PyTorch (requires Stanford login).
Michael A. Nielsen. Neural Networks and Deep Learning
Eugene Charniak. Introduction to Deep Learning

Conferences: ACL (Association for Computational Linguistics), EACL (European Association for Computational Linguistics), COLING (International Conference on Computational Linguistics), ICML (International Conference on Machine Learning), IJCNLP (International Joint Conference on Natural Language Processing), AAAI (American Association of Artificial Intelligence), ECAI (European Conference on AI), HLT/NAACl (Human language Technology/ North American Association for Computational Linguistics), ICON (International Conference on Natural Language Processing) etc.

CS551	Intro to Deep Learning	3-0-0-6	CS

This course will provide basic understanding of deep learning and how to solve classification

problems having large amount of data. In this course several public domain tools will be demonstrated to build deep learning network.

Course content will be as follows: Brief introduction of big data problem, Overview of linear algebra, probability, numerical computation Scalars, vectors, matrix, tensors, norms, eigen value, eigenvector, singular value decomposition, determinant, Probability distribution, bayes rule, conditional probability, variance, covariance, Overflow, underflow, gradient based optimization, least square,- Neural network - Perceptron, Multi-level perceptron, Universal approximation theorem,--Tutorial for Tools, Keras, Theano, Tensor flow, Demo using MNIST. Deep learning network, Shallow vs Deep network, Deep feedforward network, Gradient based learning - Cost function, soft max, sigmoid function, Hidden unit - ReLU, Logistic sigmoid, hyperbolic tangent, Architecture design, Back propagation algorithm - Chain rule of calculus, SGD, Regularization - parameter norm penalties, drop out, noise robustness, early stopping, Batch normalization, Optimization for training deep model- Adagrad, Nesterov momentum. Advanced topics: Convolutional Neural Network, Recurrent Neural Network/ Sequence modeling, Practical applications - MNIST, etc.

Texts/References:

Ian Goodfellow, YoshuaBengio and Aaron Courville, “Deep Learning”

Richard S. Sutton & Andrew G. Barto, Reinforcement Learning: An Introduction” (available online)

Jerome H. Friedman, Robert Tibshirani, and Trevor Hastie, “The elements of statistical learning”

CS563	Natural Language Processing	3-0-0-6	CS

Course Contents:

Intro to NLP

Simple Word Vector representations: word2vec, GloVe:Distributed Representations of Words and Phrases and their Compositionality, [Efficient Estimation of Word Representations in Vector Space, Advanced word vector representations: language models, GloVe: Global Vectors for Word Representation, PoS tagging and named entity recognition, Language modeling and other tasks,Opinion Mining Parsing, Sentence classification, Machine Translation, Dynamic Memory Networks, Question Answering, Natural Language Generation and Summarization, Contextual Word Representations: BERT

Text and References:

Dan Jurafsky and James H. Martin.Speech and Language Processing (3rd ed. draft)

Jacob Eisenstein.Natural Language Processing

Yoav Goldberg.A Primer on Neural Network Models for Natural Language Processing

Ian Goodfellow, YoshuaBengio, and Aaron Courville.Deep Learning

Delip Rao and Brian McMahan.Natural Language Processing with PyTorch (requires Stanford login).

Michael A. Nielsen.Neural Networks and Deep Learning

Eugene Charniak.Introduction to Deep Learning

Conferences:
ACL (Association for Computational Linguistics), EACL (European Association for Computational Linguistics), COLING (International Conference on Computational Linguistics), ICML (International Conference on Machine Learning), IJCNLP (International Joint Conference on Natural Language Processing), AAAI (American Association of Artificial Intelligence), ECAI (European Conference on AI), HLT/NAACl (Human language Technology/ North American Association for Computational Linguistics), ICON (International Conference on Natural Language Processing) etc.

EE525	Digital Image Processing	3-0-0-6	EE

Introduction to Digital Image Processing & Applications, Sampling, Quantization, Basic Relationship between Pixels, Imaging Geometry, Image Transforms, Image Enhancement, Image Restoration, Image Segmentation, Morphological Image Processing, Shape Representation and Description, Object Recognition and Image Understanding, Texture Image Analysis, Motion Picture Analysis, Image Data Compression

Texts/References:

1. Rafael C. Gonzalez and Richard E. Woods, Digital Image Processing, Pearson

2. Milan Sonka, Vaclav Hlavac and Roger Boyle, Image Processing, Analysis and Machine Vision, Springer

3. Anil K. Jain, Fundamentals of Digital Image Processing, Prentice Ha

MA564	Linear Algebra and Optimization Techniques	3-0-0-6	MA

Texts/References:

CS566	Advanced Machine Learning	3-0-0-6	CS

This course will concentrate on some advanced topics of machine learning like graphical models, auto-encoders, GANs, reinforcement learning, time series forecasting, advanced unsupervised classification algorithms, neural architectures for sequence and graph-structured predictions. When appropriate, the techniques will be linked to applications in translation, conversation modeling, bioinformatics, and information retrieval.

Prerequisites: CS 564 or an equivalent introductory course on machine learning. The course assumes basic knowledge of probability, statistics, and linear algebra.

Syllabus:

1) Mathematics of machine learning, Overview of supervised, unsupervised learning and Multi-task learning
2) Undirected graphical models: Undirected graphical models: overview, representation of probability distribution and conditional independence statement, Factorization, CRFs, Applications to NLP
3) Deep Networks for Sequence Prediction: Encoder-decoder models (case study translation), Attention models, LSTM, Memory Networks
4) Deep Network for Generation: Sequence to Sequence Models, Variational Autoencoders, Generative Adversarial Networks (GANs), Pointer Generator Networks, Transformer Networks, Learning Representations, Learning representations for text
5) Models for continuous variables: Time series forecasting
6) Reinforcement Learning: Q learning, Policy gradients, Markov Decision Process, Deep Q learning
7) Modern clustering techniques: Multi-objective optimization for clustering, Deep learning for clustering, Online Learning, Mistake Bounds
8) Recent topics for solving various problems of natural language processing, bioinformatics, information retrieval.

Books:

- Kevin P. Murphy. Machine Learning: A Probabilistic Perspective. MIT Press 2012
- Ian Goodfellow, YoshuaBengio and Aaron Courville. Deep Learning. MIT Press 2016

Other relevant textbooks:

- Yoav Goldberg. 2016. A primer on neural network models for natural language processing. *J. Artif. Int. Res.* 57, 1 (September 2016), 345-420.
- R. G. Cowell, A. P. Dawid, S. L. Lauritzen and D. J. Spiegelhalter. "Probabilistic Networks and Expert Systems". Springer-Verlag. 1999.

M. I. Jordan (ed). "Learning in Graphical Models". MIT Press. 1998

Course No.: CS695	**Name: Project I**	**Credits: 0-0-20-20**	**Prerequisites: NIL**

Syllabus:

The project can span the course Project-II. Hence it is expected that the problem specification and the milestones to be achieved in solving the problem are clearly specified. The project is encouraged to be carried out with industry.

Course No.: CS696	**Name: Project II**	**Credits: 0-0-0-24**	**Prerequisites: NIL**

Syllabus:

The students who work on a project are expected to work towards the goals and milestones set in course Project-I. At the end there would be demonstration of the solution and possible future work on the same problem. A dissertation outlining the entire problem, including a literature survey and the various results obtained along with their solutions is expected to be produced. The project is encouraged to be carried out with industry.

Basket of Electives:

Basket-1 (First Semester**):**

EE526: Digital Video Processing
MA 5xx: Linear Algebra
CS5xx: Cloud Computing
CS5xx: Data Mining Concepts

Basket-2 (First Semester**):**

CS502: Pattern Recognition
ME501: Robotics
CS 5xx: Logic in Computer Science and AI
CS5xx: Deep Learning

Basket-1 (Second Semester):
CS 5xx: Artificial Intelligence-II
CS 5xxx: Probabilistic and Approximate Reasoning
CS5xx: Edge Computing
CS5xx: Evolutionary Computation

Basket-2 (Second Semester):
MA 5xxx: Optimization Techniques
CS 5xx: Conversational AI
CS5xx: Text Mining
CS 5xx: Advanced ML
CS 5xx: Introduction to Bioinformatics

List of Electives:

CS 5xx: Artificial Intelligence-II
CS502: Pattern Recognition
CS 5xx: Conversational AI
CS5xx: Text Mining
CS5xx: Cloud Computing
CS5xx: Edge Computing
CS5xx: Evolutionary Computation
CS 5xx: Deep Learning
ME501: Robotics
EE526: Digital Video Processing
CS 5xx: Logic in Computer Science and AI
CS 5xx: Introduction to Bioinformatics
CS5xx: Data Mining Concepts
CS 5xxx: Probabilistic and Approximate Reasoning
MA 5xx: Linear Algebra
MA 5xxx: Optimization Techniques

More electives will be added in time

Syllabus of Electives:

CS 502: Pattern Recognition

Syllabus: Introduction to Pattern Recognition: Learning paradigms, Supervised and unsupervised learning; Bayesian decision theory: Minimum error rate classifier; Parameter estimation: Maximum likelihood and Bayesian Estimation; Hidden Markov models; Nonparametric techniques: Nearest neighbor rules, Parzen windows; Decision trees: Axis-parallel, Oblique, Impurity measures; Feature selection: Forward, backward search; Component analysis and discriminate functions: Principal component analysis, Fisher linear discriminate, Perceptron, Support vector machines; Generalization ability of learning methods: Bias and variance, Regularization; Bootstrapping, Boosting, Bagging; Unsupervised learning and clustering: k-Means methods.

Texts:

1. R. O. Duda, P. E. Hart and D. G. Stork, Pattern classification, John Wiley & Sons, 2002.
2. S. Theodoridis and K. Koutroumbas , Pattern Recognition, 4th Edition, Academic Press, 2008.

References:

1. C. M. Bishop, Neural Networks for Pattern Recognition, Oxford University Press, 1995.
2. V. N. Vapnik, The Nature of Statistical Learning Theory, Springer, 2000.
3. N. Cristianini and J. Shawe-Taylor, An Introduction to Support Vector Machines, Cambridge University Press, 2000.
4. Selected Research Papers.

Advanced Topics in Artificial Intelligence (AI-II)
Prerequisites: AI, Machine Learning

Introduction to the course
Knowledge Representation: Ontology, Knowledge Graph, Semantic Web
Uncertain Knowledge and Reasoning: Quantifying uncertainty, Probabilistic Reasoning, Probabilistic Reasoning over time, Multi-agent decision making
Markov Decision Processes: Policy evaluation, Policy improvement, Policy iteration, Value iteration
Reinforcement Learning: Monte Carlo, SARSA, Q-learning, Exploration/Exploitation, Function approximation, Deep reinforcement learning
Machine Learning: Clustering, Support Vector Machine, Deep Neural Networks (CNN, RNN, Auto-encoder)
Evolutionary Computation: Genetic Algorithm, Ant Colony Optimization, Particle Swarm

Optimization, Differential Evolution
Conversational AI, Explainable AI, Understanding AI Ethics and Safety

Books:
References:

1. S. Russel and P. Norvig. Artificial Intelligence: A Modern Approach (Third Edition), Prentice Hall, 2009
2. E. Rich and K. Knight, Artificial Intelligence, Addison Wesley, 1990
3. Ian Goodfellow, YoshuaBengio and Aaron Courville, Deep Learnng, MIT Press, 2016
4. Daphne Koller and Nir Friedman, Probabilistic Graphical Models: Principles and Techniques, MIT Press, 2009.
5. Sutton and Barto. Reinforcement Learning: An Introduction. Available free online.
6. Hastie, Tibshirani, and Friedman. The elements of statistical learning. Available free online.

Journals and Conference Proceedings:

Artificial Intelligence, Machine Learning, ACL Anthology, COLING, ICML, ECML, Proceedings of Uncertainty in AI, ICCV, ICLR etc.

CS5XXX:Conversational Artificial Intelligence

Thanks to the increasing rise of interest in chatbot from the industry, conversational AI (Artificial Intelligence) is a new hot research field in Natural Language Processing (NLP), Machine Learning and Deep Learning. The main goal of conversational AI is to generate a human-like conversation. However, it is a challenging task due to the complex nature of human conversations, co-reference, etc. Conversations are broadly categorized into two classes: task-oriented and chit-chat (also called as non-task oriented). Both kinds of conversations are governed by different factors or pragmatics, such as topic, interlocutors' personality, argumentation logic, viewpoint, intent, and so on. It is thus extremely important to properly model all these factors for effective conversational analysis. In this very emerging and advanced course, we will discuss the modern technologies of deep learning and word representations such as BERT, GPT-2, Seq2seq, Transformer for conversation analysis and generation. We will also take a step ahead from generation to classification by shedding light on the emotion recognition in conversation, emotion-oriented dialogue generation, conversational question-answering (QA), intent classification etc.This course will help students, and researchers to understand both the basic and advanced techniques of conversational AI.

Syllabus:

Introduction and basic Concepts to Dialogue Systems.

Introduction to Machine Learning, Natural Language Processing and Deep Learning

Recurrent Neural Networks, LSTM, Bi-LSTM, Encoder-Decoder, Attention

Natural Language Understanding: Dialogue Act, Intent detection and Slot filling

Dialogue Management: Sequence Models and Reinforcement Learning, Reinforcement Learning for Dialogue Management

Natural Language Generation:Language Models for generation, Word Embedding, Word2Vec, Transformer network, BERT and GPT-2 language modeling for generation

Task oriented dialogue generation using advanced technique

Conversational Question-Answering (QA), Multimodal Fusion Techniques, Multimodal dialogue generation

Emotion recognition and sentiment analysis in conversations, Affective and emotion-oriented dialogue generation

KG driven NLG and Personalization

Knowledge graph aware Natural Language Generation, Personalized Dialogue Generation

Conversation Machine Translation

Texts and References

Spoken Language Understanding: Systems for Extracting Semantic Information from Speech by Gokhan Tur and Renato De Mori

Speech and Language Processing by Dan Jurafsky and James H. Martin

- Chapter 29: Dialog Systems and Chatbots
- Chapter 30: Advanced Dialog Systems

Spoken Dialogue Systems by Kristiina Jokinen and Michael McTear, also available at Amazon.

Conferences and Journals: ACL, EMNLP, COLING, NAACL, EACL, IJCNLP, AACL, Computational Linguistics, Transaction on ACL, IEEE Transaction on Affective Computing, ACM Transaction on Intelligent System, ACM Transaction on Human Computer Interaction

CS5XX: Text Mining

Given the dominance of text information over the Internet, mining high-quality information from text becomes increasingly critical. The actionable knowledge extracted from text data facilitates our life in a broad spectrum of areas, including business intelligence, information acquisition, social behavior analysis and decision making. In this course, we will cover important topics in text mining including: basic natural language processing techniques, document representation, text categorization and clustering, document summarization, sentiment analysis, probabilistic topic models.

Introduction: Fundamental concepts Natural language processing: Part-of-speech tagging, chunking, syntax parsing and named entity recognition.

Document representation: Vector Space Model

Text categorization: Basic supervised text categorization algorithms, including Naive Bayes, k Nearest Neighbor (kNN) and Logistic Regression.

Text clustering: Two typical types of clustering algorithms, i.e., connectivity-based clustering (a.k.a., hierarchical clustering) and centroid-based clustering (e.g., k-mean

Topic modeling: General idea of topic modeling, two basic topic models, i.e., Probabilistic Latent Semantic Indexing (pLSI) and Latent Dirichlet Allocation (LDA), and their variants for different application scenarios, including classification, collaborative filtering, and hierarchical topical structure modeling.

Document summarization: Extractive and Abstractive summarization

Sentiment Analysis: Coarse-grained and Fine-grained analysis, Machine Learning for Sentiment Analysis

Text and Reference books:

Mining Text Data. Charu C. Aggarwal and ChengXiangZhai, Springer, 2012.
Speech & Language Processing. Dan Jurafsky and James H Martin, Pearson Education India, 2000.
Introduction to Information Retrieval. Christopher D. Manning, Prabhakar Raghavan, and HinrichSchuetze, Cambridge University Press, 2007.
Conferences and Journals: ICML, KDD, ICLR, ACL, EMNLP, COLING, EACL, NAACL, IEEE TKDE, ACM TKDD, ACM Transaction on Low-resources

CS 5XX: Logic in Computer Science and Artificial Intelligence

Syllabus:

Introduction: definition, history, Logic and the foundations of mathematics, Logic in computer science

Propositional Logic: Syntax- Alphabet, Well-formed formulas, Unique readability,

Propositional Logic: Semantics- Interpretations and satisfaction, Validity, Satisfiability, and unsatisfiability, Logical implication, Substitution.

Propositional Logic: Decision Procedures-Truth tables, Semantic trees, Complexity (NP-completeness), Horn clauses.

Propositional Logic, A Formal System- Axioms, inference rules, and deductions, Monotonicity, Soundness, Deduction Theorem, Consistency and inconsistency, Completeness, Compactness

Propositional Logic- Resolution: Literals and clauses, Resolution, Soundness, Completeness, Complexity

First-Order Logic- Syntax and Semantics: Terms and formulas, Structures, assignments, and semantics, Satisfiability, validity, and logical implication.

First-Order Logic- Definability: Preservation under isomorphism; Elementary equivalence vs. isomorphism; Classes of structures, relations, and queries; Explicit and implicit definability; Compactness and applications to definability.

First-Order Logic- Normal Forms: Prenex normal form, Skolem normal form, Elimination of function symbols, Elimination of equality.

First-Order Logic- A Formal System;Propositional reasoning, The Deduction Theorem, Substitutions, Axioms, Generalization.

First-Order Logic- Completeness and Consequences: Completeness, Lowenheim-Skolem Theorem, Compactness Theorem.

First-Order Logic- Undecidability and Incompleteness: Recursive enumerability of validity, Non-recursive enumerability of satisfiability, Decidable quantational classes, Godel'sIncompleteness Theorem

Recommended Readings:

M. Gardner: Logic Machines and Diagrams, 1982
A. Feferman: Politics, Logic, and Love: the life of Jean Van Heijenoort, 1993

M. Davis: The Universal Computer: The Road from Leibniz to Turing , 2011.

CS5XX: Evolutionary Computation

Introduction to Evolutionary Computation: Biological and artificial evolution, Evolutionary computation and AI, Different historial branches of EC, e.g., GAs, EP, ES, GP, etc., A simple evolutionary algorithm

Search Operators: Recombination/Crossover for strings (e.g., binary strings), e.g., one-point, multi-point, and uniform crossover operators, Mutation for strings, e.g., bit-flipping, Recombination/Crossover and mutation rates, Recombination for real-valued representations, e.g., discrete and intermediate recombinations, Mutation for real-valued representations, e.g., Gaussian and Cauchy mutations, self-adaptive mutations, etc., Why and how a recombination or mutation operator works

Selection Schemes: Fitness proportional selection and fitness scaling, Ranking, including linear, power, exponential and other ranking methods, Tournament selection, Selection pressure and its impact on evolutionary search

Search Operators and Representations: Mixing different search operators, An anomaly of self-adaptive mutations, The importance of representation, e.g., binary vs. Gray coding, Adaptive representations

Evolutionary Combinatorial Optimisation: Evolutionary algorithms for TSPs, Evolutionary algorithms for lecture room assignment, Hybrid evolutionary and local **search algorithms**

Coevolution: Cooperative co-evolution, Competitive coevolution

Niching and Speciation: Fitness sharing (explicit and implicit), Crowding and mating restriction

Constraint Handling: Common techniques, e.g., penalty methods, repair methods, etc.

Analysis, Some examples

Genetic Programming: Trees as individuals, Major steps of genetic programming, e.g., functional and terminal sets, initialisation, crossover, mutation, fitness evaluation, etc., Search operators on trees, Automatically defined functions, Issues in genetic programming, e.g., bloat, scalability, etc., Examples

Multiobjective Evolutionary Optimisation: Pareto optimality, Multiobjective evolutionary algorithms

Learning Classifier Systems: Basic ideas and motivations, Main components and the main cycle, Credit assignment and two approaches

Theoretical Analysis of Evolutionary Algorithms: Schema theorems, Convergence of EAs, Computational time complexity of EAs, No free lunch theorem

Recommended Books:

1. Handbook on Evolutionary Computation, T. Baeck, D. B. Fogel, and Z. Michalewicz (eds.), IOP Press, 1997.
2. Genetic Algorithms in Search, Optimisation& Machine Learning, D E Goldberg, Addison-Wesley, 1989
3. Multi-Objective Optimization using Evolutionary Algorithms, by Kalyanmoy Deb, Wiley

CS5XX: Data Mining Concepts

Introduction to Data Mining: Data Mining Goals, Stages of the Data Mining Process, Data Mining Techniques, Knowledge Representation Methods, Applications

Data preprocessing: Data cleaning, Data transformation, Data reduction, Discretization and generating concept hierarchies

Data mining knowledge representation: Task relevant data, Background knowledge, Interestingness measures, Representing input data and output knowledge, Visualization techniques

Attribute-oriented analysis: Attribute generalization, Attribute relevance, Class comparison, Statistical measures, Experiments with Weka - using filters and statistics

Data mining algorithms: Association rules: Motivation and terminology, Basic idea: item sets, Generating itemsets and rules efficiently, Correlation analysis

Data mining algorithms: Classification- Basic learning/mining tasks, inferring rudimentary rules: 1R algorithm, Decision trees, Covering rules

Data mining algorithms: Prediction - The prediction task, Statistical (Bayesian) classification, Bayesian networks, Instance-based methods (nearest neighbor), Linear models, Experiments with Weka - Prediction

Mining real data: Preprocessing data from a real medical domain (310 patients with Hepatitis C)., Applying various data mining techniques to create a comprehensive and accurate model of the data**.**

Clustering: Basic issues in clustering, First conceptual clustering system: Cluster/2, Partitioning methods: k-means, expectation maximization (EM), Hierarchical methods: distance-based agglomerative and divisible clustering, Conceptual clustering.

Advanced techniques, Data Mining software and applications: Text mining: extracting attributes (keywords), structural approaches (parsing, soft parsing), Bayesian approach to classifying text, Web mining: classifying web pages, extracting knowledge from the web.

Recommended Readings

1. Ian H. Witten and Eibe Frank, *Data Mining: Practical Machine Learning Tools and Techniques (Second Edition)*, Morgan Kaufmann, 2005, ISBN: 0-12-088407-0.

CS 5XX: Introduction to Bioinformatics

Introduction to Bioinformatics and Key Online Bioinformatics Resources: NCBI & EBI Biology is an information science, History of Bioinformatics, Types of data

Application areas: Introduction to upcoming segments, NCBI & EBI resources for the molecular domain of bioinformatics, Focus on GenBank, UniProt, Entrez and Gene Ontology.

Sequence Alignment: DNA and Protein Database Searching Homology, Sequence similarity, Local and global alignment, Database searching with BLAST.

Advanced Database Searching: PSI-BLAST, Profiles and HMMs, Protein structure comparisons.
Structural Bioinformatics: Protein structure function relationships, Protein structure and visualization resources, Structural genomics, Homology modeling, Inferring protein function

from structure.

Structure Based Drug Discovery and Biomolecular Simulations: Small molecule docking methods, Protein motion and conformational variants, Bioinformatics in drug discovery.

Genome Informatics and High Throughput Sequencing: Searching genes and gene functions, Genome databases, Variation in the Genome, Highthroughput sequencing technologies, biological applications, bioinformatics analysis methods.

Genes and Disease Human examples: Genomics and human health, The promise and potential of shifting medicine from a reactive practice of treating symptoms and diseases, to one where disease risk is diagnosed early or even managed prior to onset.

Proteomics and the Transcriptome: Processing and extracting biological information from proteomic and transcriptomic datasets, Analysis of RNA-Seq data, Differential expression tests, Avoiding P-value misuse, Hands-on analysis of RNA-Seq data.

Systems Biology: From genome to phenotypes. Integration of genome-wide data sets into their functional context, Analysis of protein-protein interactions, pathways and networks, Modeling and simulation of systems and networks, Computational methods of network modeling

Suggested Readings

1. **BIOINFORMATICS ALGORITHMS Hardcover – January 1, 2018 by Phillip Compeau, Pavel Pevzner**
2. **Bioinformatics and Functional Genomics 3rd Edition,**by Jonathan Pevsner
3. Journals: Bioinformatics, BMC Bioinformatics; and Conferences

ME501 Robotics: Advanced Concepts and Analysis (3-0-0-6) Prerequisite NIL

Introduction to robotics: brief history, types, classification and usage and the science and technology of robots. Kinematics of robot: direct and inverse kinematics problems and workspace, inverse kinematics solution for the general 6R manipulator, redundant and over-constrained manipulators. Velocity and static analysis of manipulators: Linear and angular velocity, Jacobian of manipulators, singularity, static analysis. Dynamics of manipulators: formulation of equations of motion, recursive dynamics, and generation of symbolic equations of motion by computer simulations of robots using software and commercially available packages. Planning and control: Trajectory planning, position control, force control, hybrid control Industrial and medical robotics: application in manufacturing processes, e.g. casting, welding, painting, machining, heat treatment and nuclear power stations, etc; medical robots: image guided surgical robots, radiotherapy, cancer treatment, etc; Advanced topics in robotics: Modelling and control of flexible manipulators, wheeled mobile robots, bipeds, etc. Future of robotics.

Texts and Reference Books

1. M. P. Groover, M. Weiss, R. N. Nagel and N. G. Odrey, “Industrial Robotics-Technology, Programming and Applications”, McGraw-Hill Book and Company (1986).
2. S. K. Saha, “Introduction to Robotics”, Tata McGraw-Hill Publishing Company Ltd. (2008).
3. S. B. Niku, “Introduction to Robotics–Analysis Systems, Applications”, Pearson Education (2001).
4. A. Ghosal, Robotics: “Fundamental Concepts and Analysis”, Oxford University Press (2008).

5. Pires, “Industrial Robot Programming–Building Application for the Factories of the Future”, Springer (2007).
6. Peters, “Image Guided Interventions – Technology and Applications”, Springer (2008).
7. K. S. Fu, R. C. Gonzalez and C.S.G. Lee, “ROBOTICS: Control, Sensing, Vision and Intelligence”, McGraw-Hill (1987).
8. J. J. Craig, “Introduction to Robotics: Mechanics and Control”, 2nd edition, Addison-Wesley **(1989).**

CS 5XX: Probabilistic and Approximate Reasoning

Overview of Probability Theory, Bayes Networks, Independence, I-Maps, Undirected Graphical Models, Bayes Networks and Markov Networks, Local Models, Template Based Representations, Exact Inference: Variable Elimination; Clique Trees, Belief Propagation, Tree Construction, Intro to Optimization, Approximate Inference: Sampling, Markov Chains, MAP Inference, Inference in Temporal Models, Learning Graphical Models : Intro, Parameter Estimation, Bayesian Networks and Shared Parameters, Structure Learning and search, Partially Observed Data, Gradient Descent, EM, Hidden Variables, Undirected Models, Causality, Utility Functions, Decision Problems, Expected Utility, Value of Information

Suggested Reading

1. Probabilistic Graphical Models, by Daphne Koller and Nir Friedman, MIT Press, 2009.
2. Journals and conferences in AI

EE526 Digital Video Processing

Representation of digital video: Introduction and fundamentals; Time-varying image formation models: Motion models, Geometric image formation; Spatio-temporal sampling: Sampling of analog and digital video, Two-dimensional rectangular and periodic sampling, Sampling of 3-D structures, Reconstruction from samples; Sampling structure conversion: Sampling rate change, Sampling lattice conversion; Two-Dimensional Motion Estimation: Optical flow based methods, Block-based methods, Pel-recursive methods, Bayesian methods based on Gibbs Random Fields; Image Compression: Lossless compression, DPCM, Transform coding, JPEG, Vector Quantization, Sub-band Coding; Video compression: Inter-frame compression methods (3-d waveform and motion-compensated waveform coding), Video compression standards (H.26X and MPEG-X); Applications of video processing: Video Indexing, Summarization, Browsing and Retrieval, Video Surveillance

Texts:

1. A. M. Tekalp, “Digital Video Processing”, Prentice Hall.
References.
2. R. C. Gonzalez, and R. E. Woods, “Digital Image Processing”, Addison-Wesley.
3. Dudgeon &Mersereau, “Multi-dimensional Digital Signal Processing”, Prentice Hall.
4. C. Poynton, “A Technical Introduction to Digital Video”, Wiley.
5. Y. Wang, J. Ostermann, and Y. Zhang, “Video Processing and Communications”, Prentice Hall.
6. K. Castleman, “Digital Image Processing”, Prentice Hall.
7. S. Mitra, “Digital Signal Processing”, 2nd Edition, McGraw Hill.

MA5XX: Optimization Techniques

Introduction to linear and non-linear programming. Problem formulation. Geo- metrical aspects of LPP, graphical solution. Linear programming in standard form, simplex, Big M and Two Phase Methods. Revised simplex method, special cases of

LP. Duality theory, dual simplex method. Sensitivity analysis of LP problem. Transportation, assignment and traveling salesman problem.

Integer programming Problems-Branch and bound method, Gomory cutting plane method for all integer and for mixed integer LP.

Unconstrained Optimization, basic descent methods, conjugate direction and Newton's methods. Acquaintance to Optimization softwares like TORA.

Books:

Hamdy A. Taha, Operations Research: An Introduction, Eighth edition, PHI, New Delhi (2007).

S. Chandra, Jayadeva, AparnaMehra, Numerical Optimization with Applications, Narosa Publishing House (2009).

A. Ravindran, Phillips, Solberg, Operation Research, John Wiley and Sons, New York (2005).

M. S. Bazaraa, J. J. Jarvis and H. D. Sherali, Linear Programming and Network Flows, 3rd Edition, Wiley (2004).

References:

D. G. Luenberger, Linear and Nonlinear Programming, 2nd Edition, Kluwer, 2003. S. A. Zenios (editor), Financial Optimization, Cambridge University Press (2002).

F. S. Hiller, G. J. Lieberman, Introduction to Operations Research, Eighth edition, McGraw Hill (2006).

MA5XX:Linear Algebra

Systems of linear equations, Matrices, Elementary row operations, Row-reduced echelon matrices, Vector spaces, Subspaces, Bases and dimension, Ordered bases and coordinates. Linear transformations, Rank-nullity theorem, Algebra of linear transformations, Isomorphism, Matrix representation, Linear functionals, Annihilator, Double dual, Transpose of a linear transformation. Characteristic values and characteristic vectors of linear transformations, Diagonalizability, Minimal polynomial of a linear transformation, Cayley-Hamilton theorem, Invariant subspaces, Direct-sum decompositions, Invariant direct sums, The primary decomposition theorem, Cyclic subspaces and annihilators, Cyclic decomposition, Rational, Jordan forms. Inner product spaces, Orthonormal bases, Gram-Schmidt process.

Suggested Readings:

1. K.Hoffman and R.Kunze, Linear Algebra, 2nd Edition, Prentice- Hall of India, 2005.
2. M.Artin,Algebra,Prentice-Hall of India, 2005
3. S.Axler, Linear Algebra Done Right, 2nd Edition, John-Wiley, 1999.
4. S. Lang, Linear Algebra, Springer UTM, 1997.

5. S.Kumaresan, Linear Algebra:A Geometric Approach,Prentice-Hall of India, 2004.