FIRST SEMESTER

MAT 5135 MATHEMATICAL FOUNDATIONS OF COMPUTATIONAL INTELLIGENCE [3 1 0 4]

Abstract

Linear Algebra and Optimization; Optimization for Machine Learning; Analytic Geometry; Linear Transformations and Linear Systems; Eigenvectors and Diagonalizable Matrices; Optimization Basics; Advanced Optimization Solutions; Constrained Optimization and Duality; Matrix Factorization; Probability and Distributions

SDL component: Topics from Constraint optimization and advanced optimization.

Course Outcomes

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

1. Apply the concepts of reasoning calculus to solve complex engineering problem
2. Apply the concepts of linear algebra to solve real world problems.
3. Apply different optimization techniques to solve complex engineering problem.
4. Apply the concept of probability and distributions to solve real world problems.

References

1. Michael Huth and Mark Ryan, *Logic in Computer Science*, 2nd Edition, Cambridge University Press, 2007.
2. Dana Richards, and Henry Hamburger, *Logic and Language Models for Computer Science*, 4th Edition, World Scientific, 2023
3. Charu C Agarwal, *Linear Algebra and Optimization for Machine Learning*, Springer, 2020
4. Marc Peter Deisenroth, A.Also Faisal, and Cheng Soon Ong, *Mathematics for Machine Learning*, Cambridge University Press, 2020
5. Kaare Brandt Petersen, and Michael Syskind Pedersen, *The Matrix Cookbook*, Technical University of Denmark, 2012

ICT 5113 GENERATIVE ARTIFICIAL INTELLIGENCE [4 0 0 4]

Abstract:

Introduction; Generative Modeling, Deep Learning, Variational Encoders, Generative Adversarial Networks, Autoregressive Models, Normalizing Flow Models, Energy-Based Models, Diffusion Models, Transformers, Music Generation, World Models, Multimodal Models, Ethical Issues in Generative AI

Self-Directed Learning: Basics of neural networks, Different types of activation functions

Course Outcomes:

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

1. Understand the principles and methods of generative AI.
2. Learn how to use popular frameworks and tools such as PyTorch, and Hugging Face.
3. Implement and customize generative models for various tasks and domains.
4. Evaluate and compare the performance and quality of generative models.
5. Explore the opportunities and risks of generative AI for individuals, businesses, and society.

References:

1. Jakub Langr, and Vladimir Bok, *GANs in Action: Deep Learning with Generative Adversarial Networks*, Manning, 2019.

2.David Foster, *Generative Deep Learning: Teaching Machines to Paint, Write, Compose and Play*, 2^{nd} Edition, O'Reilly, 2023.

3.Joseph Babcock, and Raghav Bali, *Generative AI with Python and TensorFlow 2: Create images, text, and music with VAEs, GANs, LSTMs, Transformer models*, Packt, 2021.

4.Altaf Rehmani, *Generative AI for Everyone*, Bluerose Publishers Pvt. Ltd, 2024

5.Oliver Caelen and Marie-Alice Blete, *Developing Apps with GPT-4 and ChatGPT,* O'Reilly, 2023

ICT 5130 THEORETICAL FONDATIONS OF MACHINE LEARNING [4 0 0 4]

Abstract

Introduction to Machine Learning, Mathematical Preliminaries, Supervised Learning-LMS, logistic regression, GDA, Naive Bayes, Kernel Methods, SVM, Model and feature selection, Deep Learning basics, Generalization and regularization, Unsupervised learning-clustering, k-means, Gaussian mixture, factor analysis, PCA, ICA, Reinforcement learning-MDPs, Bellman equations, value and policy iteration, LQR, LQG, Q-learning, policy search, POMDPs,

Course Outcomes

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

1. Articulate machine learning algorithms in data-driven knowledge discovery
2. Identify the suitability of discriminative or generative supervised model for a given problem
3. Choose an appropriate unsupervised model for a given learning problem
4. Formulate an appropriate reinforcement model for a given learning scenario
5. Apply diagnostics for debugging learning algorithms

References

1. Kevin P Murphy, *Machine Learning: A Probabilistic Perspective*, MIT Press, 2012.
2. Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar., *Foundations of Machine Learning*, MIT Press, 2012.
3. Christopher M.Bishop., *Pattern Recognition and Machine Learning (2e)*, Springer, 2013.
4. Charu C Aggarwal, Neural Networks and Deep Learning, 2nd Edition, Springer, 2023
5. Richard S.Sutton and Andrew G.Barto, *Reinforcement Learning*, 2nd Edition, MIT Press, 2018

ICT 5125 PROBABILISTIC GRAPHICAL MODELS [4 0 0 4]

Abstract

Introduction; Mathematical Preliminaries; Representation: Bayesian Networks Representation, Undirected Graphical Models, Gaussian Network Models; Inference: Variable Elimination, Clique Trees, Inference as Optimization, Particle-Based Approximate Inference; Learning: Learning Graphical Models, Parameter Estimation, Structure Learning in Bayesian Networks, Partially Observed Data, Learning Undirected Models, Real-world applications.

Self-Directed Learning: Probability distributions, Basic concepts in probability, Random variables and joint distributions, Independence and conditional independence

Course Outcomes

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

1. Demonstrate the understanding of mathematical framework of probabilistic graphical models
2. Implement the basic algorithms for probabilistic inference in graphical models
3. Implement the basic algorithms for learning graphical models
4. Identify and applying Bayesian principles behind modeling domain knowledge under uncertainty
5. Recognize the difference between statistical and causal models

References

1. Daphne Koller and Nir Friedman, *Probabilistic Graphical Models: Principles and Techniques*, MIT Press, 2009
2. Luis Enrique Sucar, *Probabilistic Graphical Models: Principles and Applications*, Springer, 2020
3. Kevin P Murphy, *Machine Learning: A Probabilistic Perspective*, MIT Press, 2012
4. David Barber, *Bayesian Reasoning and Machine Learning*, Cambridge University Press, 2012
5. Adnan Darwiche, *Modeling and Reasoning with Bayesian networks*, Cambridge University Press, 2014

ICT 5126 FAIRNESS AND MACHINE LEARNING [3 0 0 3]

Abstract

Introduction; Classification: Supervised learning, Sensitive characteristics, Formal non-discrimination criteria, Calibration and sufficiency, Relationships between criteria; Causal Models: Simpson's paradox, Graphs, Structural causal models; Causality; Effect of Interventions: Adjustment Formula, Backdoor Criterion, Front-Door Criterion; Counterfactuals and applications; Traditional tests for discrimination; Testing discrimination in algorithmic systems; Datasets; Fairness mechanism; Legal and policy perspective

Self-Directed Learning: Basics of algorithmic fairness and explainability

Course Outcomes

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

1. Differentiate between black box models and the explainable models
2. Make use of tools for model visualization, and make interpretation
3. Make model interpretation for a given learning problem
4. Explain deep learning models
5. Apply fairness mechanisms to learning algorithms.

References

1. Solon Barocas, Moritz Hardt and Arvind Narayanan, *Fairness and Machine Learning*, failml.org, 2023
2. Judea Pearl, *Causality: Models, Reasoning, and Interference*, 2nd Edition, Cambridge University Press, 2009
3. Jonas Peters, Dominik Janzing, and Bernhard Scholkopf, *Elements of Causal Inference: Foundations and Learning Algorithms*, MIT Press, 2017
4. Madelyn Glymour, Judea Pearl, and Nicholas P. Jewell, *Causal Inference in Statistics: A Primer*, 1st Edition, Wiley, 2016
5. Serg Masis, *Interpretable Machine Learning with Python*, Packt Publishing Ltd

HUM 5071 RESEARCH METHODOLOGY AND TECHNICAL COMMUNICATION

[1 0 3 -]

Abstract

Mechanics of Research Methodology: Basic concepts: Types of research, Significance of research, Research framework, Case study method, Experimental method, Sources of data, Data collection using questionnaire, Interviewing, and experimentation. Research formulation: Components, selection and formulation of a research problem, Objectives of formulation, and Criteria of a good research problem. Research hypothesis: Criterion for hypothesis construction, Nature of hypothesis, need for having a working hypothesis, Characteristics and Types of hypothesis, Procedure for hypothesis testing, Sampling methods- Introduction to various sampling methods and their applications. Data Analysis: Sources of data, Collection of data, Measurement and scaling technique, and Different techniques of Data analysis. Thesis Writing and Journal Publication: thesis writing, journal and conference papers writing, IEEE and Harvard styles of referencing, Effective Presentation, Copyrights, and avoiding plagiarism.

Self-Directed Learning: IEEE and Harvard styles of referencing

References

1. Ranjit Kumar, Research Methodology: *A Step-by-Step Guide for Beginners*, SAGE, 2005.
2. Geoffrey R. Marczyk, David DeMatteo & David Festinger, *Essentials of Research Design and Methodology*, John Wiley & Sons, 2004.
3. John W. Creswel , *Research Design: Qualitative, Quantitative, and Mixed Methods Approaches*, SAGE, 2004
4. Suresh C. Sinha and Anil K. Dhiman, *Research Methodology (2 Vols-Set)*, Vedam Books, 2006.
5. C. R. Kothari, *Research Methodology: Methods and Techniques*, New Age International Publisher, 2008.
6. Donald R Cooper & Pamela S Schindler, *Business Research Methods*, McGraw Hill International, 2007.
7. R. Pannershelvam, *Research Methodology*, Prentice Hall, India, 2006
8. Manfred Max Bergman, *Mixed Methods Research*, SAGE Books, 2006.
9. Paul S. Gray, John B. Williamson, David A. Karp, John R. Dalphin, *The Research Imagination*, Cambridge University press, 2007.
10. Cochrain & Cox, *Experimental Designs*, II Edn. Wiley Publishers, 2006.

ICT 5143 ALGORITHMIC FAIRNESS AND GRAPHICAL MODELS LAB [0 0 3 2]

Abstract

Students will be able to experiment on various aspects of explainability, and algorithm fairness. For explainability they should be able to use causality theory tools in terms of graphical models. Students will be using tools like BRML Toolbox, and PGMPy.

Self-Directed Learning: Basics of tools like BRML toolbox and PGMPy

Course Outcomes

At the end of this course, students will be able to

1. Define the problem statement explicitly specifying algorithmic fairness
2. Demonstrate the methodology followed for quantifying algorithmic fairness and explainability
3. Demonstrate the usage of graphical models for causality studies in explainability problems.

References

1. Solon Barocas, Moritz Hardt and Arvind Narayanan, *Fairness and Machine Learning*, failml.org, 2021
2. Judea Pearl, *Causality: Models, Reasoning, and Interference*, 2nd Edition, Cambridge University Press, 2009

3. Jonas Peters, Dominik Janzing, and Bernhard Scholkopf, *Elements of Causal Inference: Foundations and Learning Algorithms*, MIT Press, 2017
4. Madelyn Glymour, Judea Pearl, and Nicholas P. Jewell, *Causal Inference in Statistics: A Primer*, 1st Edition, Wiley, 2016
5. Serg Masis, *Interpretable Machine Learning with Python*, Packt Publishing Ltd
6. Luis Enrique Sucar, *Probabilistic Graphical Models: Principles and Applications*, Springer, 2020

ICT 5144 MACHINE INTELLIGENCE CAPSTONE PROJECT [0 0 3 2]

Abstract

Capstone project in decision science will contributes to the development of machine intelligence. The projects will be across three research axes: machine learning, deep learning, and probabilistic graphic models. The duration of capstone project will be of 12 weeks. As a part of this, student will submit detailed problem statement explicitly specifying the objectives (3^{rd} week). A mid-term evaluation of the capstone project work will be done in 6^{th} week. Students will submit the synopsis in 11^{th} week. The final evaluation will be conducted in 12^{th} week. The evaluation is based on demo and final report submission followed by viva-voce.

Self-Directed Learning: Depending on the domain students will learn to use tools like, TensorFlow, Scikit Learn, and PyTorch

Course Outcomes

At the end of this course, students will be able to

1. Define the problem statement explicitly specifying objectives.
2. Demonstrate the methodology followed for the given problem statement
3. Implementation of the capstone project.

References

1. Chip Huyen, *Designing Machine Learning Systems*, O'Reilly Media, Inc, 2022
2. Stuart Russell, and Peter Norvig, *Artificial Intelligence A Modern Approach*, 4^{th} Edition, Pearson, 2021.
3. Ian Goodfellow and Yoshua Bengio and Aaron Courville, *Deep Learning*, MIT Press, 2016
4. Christopher M. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2016.
5. Francois Chollet, *Deep Learning with Python*, Manning, 2017
6. Luis Enrique Sucar, *Probabilistic Graphical Models: Principles and Applications*, Springer, 2020

SECOND SEMESTERS

ICT 5215 CONVEX OPTIMIZATION AND APPLICATIONS [3 1 0 4]

Abstract
Introduction, Convex Sets: Convexity preserving operations, Separating and supporting hyperplanes, Dual cones; Convex Functions: Basic properties, Conjugate functions, Quasiconvex functions; Convex Optimization Problems: Optimization problems, Linear and quadratic convex optimization problems, Vector optimization; Duality: Lagrange dual function and problem; Applications: Approximation and Fitting, Statistical Estimation

Self-Directed Learning: Basics of CVXPY and CPLEX toolbox

Course Outcomes
At the end of this course, students will be able to:

1. Apply the concepts of mathematical optimization, convex sets and convex function and to solve complex engineering problem.
2. Solve convex optimization problems using different techniques learnt.
3. Apply the concept of duality to solve the optimization problem.
4. Apply different techniques learnt in approximation and fitting to solve the real world problem.
5. Solve an optimization problem using CVXPY and CPLEX

References

1. Stephen Boyd and Lieven Vandenberghe, *Convex Optimization*, Cambridge University Press, 2016
2. Dimitri P.Bertsekas, Angelia Nedic, and Asuman E.Ozdaglar, *Convex analysis and Optimization*, Athena Scientific, 2003
3. Dmitri P.Bertsekas, *Convex Optimization Algorithms*, Athena Scientific, 2015
4. Roger Fletcher, *Practical Methods of Optimization*, 2nd Edition, Wiley, 2000

ICT 5216 DATA MINING AND KNOWLEDGE DISCOVERY [4 0 0 4]

Abstract

Introduction to Relational Database Design, Data Warehousing, OLAP operations, Warehouse schema, Data Warehousing Architecture, Data Warehouse Backend Process. Data Pre-processing, Data cube, Sampling, Discretization and concept hierarchy generation, Segmentation by natural partitioning. Introduction to Data mining, Association rules mining, market based analysis, Apriori Algorithm, Partition Algorithm, Pincer – Search Algorithm, Dynamic item set counting algorithm, FP-tree growth Algorithm,PC Tree, Multilevel association rules, Approaches to mining multilevel association rules, correlation analysis, Issues and challenges in Data mining. Introduction to Clustering Techniques, Clustering paradigms, Partitioning Algorithms, k – Medoid & k- means Algorithms, Introduction to Classification and Prediction, Tree Construction principle, Decision Tree Construction

Algorithm, Data Partitioning Techniques, Interquery, Intraquery, Intraoperation and Interoperation Parallelisms, Distributed data storage, Distributed Transactions, Commit Protocols, Data Lake

Self-Directed Learning: Association rules mining, market based analysis, Apriori Algorithm, Partition Algorithm

Course Outcomes

By the end of this course the student should be able to

1. Learn the concepts of Data Preprocessing.
2. Understand Data Warehouse architecture, OLAP operations.
3. Analyze association rule mining algorithms to discover frequent itemsets.
4. Understand Clustering techniques, build Decision Tree for classification, Parallel Databases, Distributed Databases and Data Lake

References

1. Jiawei Han and Micheline Kamber, *Data Mining Concepts and Techniques*, 3rd Edition, Morgan Kauffmann Publishers, 2012
2. Manoj Kukreja, and Danil Zburivsky, *Data Engineering with Apache Spark, Delta Lake, and Lakehouse*, Packt Publication, 2021
3. Arun K Pujari, *Data Mining Techniques*, 4th Edition, Universities Press India, 2016
4. Silberschatz, Korth and Sudarshan, *Database System Concepts*, 7th Edition, McGraw Hill, 2021

PROGRAM ELECTIVE I **[3 0 0 3]**

PROGRAM ELECTIVE II **[3 0 0 3]**

PROGRAM ELECTIVE III **[4 0 0 4]**

OPEN ELECTIVE **[3 0 0 3]**

ICT 5245 KNOWLEDGE ENGINEERING LAB [0 0 3 2]

Abstract

Pre-processing the raw-datasets using data mining software/tool, Applying data mining techniques such as association rule mining, clustering, classification on the pre-processed data using software tools, , Implementation of data mining algorithms based on association rule mining, clustering, classification using python, Design and development of Mini Project.

Course Outcomes

At the end of this course, students will be able to

1. Use appropriate data mining software/tool for Pre-processing
2. Apply association rule mining technique on the pre-processed data
3. Apply clustering, classification technique on the pre-processed data
4. Design and development of Mini Project

References

1. Gheorghe Tecuci, Dorin Marcu, Mihai Boicu, and David A Schum, *Knowledge Engineering*, 1st Edition, Cambridge University Press, 2016
2. Hamed Fazlollahtabar, *Knowledge Engineering: The Process Paradigm*, 1st Edition, CRC, 2020

ICT5246 DECISION SCIENCE CAPSTONE PROJECT [0 0 3 2]

Abstract

Capstone project in decision science will contributes to the development of digital intelligence. The projects will be across four research axes. The duration of capstone project will be of 12 weeks. As a part of this, student will submit detailed problem statement explicitly specifying the objectives (3rd week). A mid-term evaluation of the capstone project work will be done in 6th week. Students will submit the synopsis in 11th week. The final evaluation will be conducted in 12th week. The evaluation is based on demo and final report submission followed by viva-voce.

Course Outcomes

At the end of this course, students will be able to

1. Define the problem statement explicitly specifying objectives.
2. Demonstrate the methodology followed for the given problem statement
3. Implementation of the capstone project.

References

1. Stephen Boyd and Lieven Vandenberghe, *Convex Optimization*, Cambridge University Press, 2016
2. Dimitri P.Bertsekas, Angelia Nedic, and Asuman E.Ozdaglar, *Convex analysis and Optimization*, Athena Scientific, 2003
3. Dmitri P.Bertsekas, *Convex Optimization Algorithms*, Athena Scientific, 2015
4. Roger Fletcher, *Practical Methods of Optimization,* 2nd Edition, Wiley, 2000

PROGRAM ELECTIVES

ICT 5411 ALGORITHMIC GAME THEORY [3 0 0 3]

Abstract

Introduction, Mechanism Design Basics, Myerson's Lemma, Algorithmic Mechanism Design, Revenue-Maximizing Auctions, Simple Near-Optimal Auctions, Multi-Parameter Mechanism Design, Spectrum Auction, Mechanism Design with Payment Constraints, Stable Matching, Selfish Routing and the Price of Anarchy, Over-Provisioning and Atomic Selfish Routing, Equilibria, Robust Price-of-Anarchy Bounds in Smooth Games, Best-Case and Strong Nash Equilibria, Best Response Dynamics, No-Regret Dynamics, Swap Regret and the Minimax Theorem, Pure Nash Equilibria and PLS-Completeness, Mixed Nash Equilibria and PPAD-Completeness

Self-Directed Learning: Concepts in non-cooperative and cooperative game theory

Course Outcomes

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

1. Understand various models of games and how they arise in various applications in computer science
2. Develop an understanding of algorithms used to solve such games
3. Model various scenarios as strategic games, and devise algorithms to solve them
4. Develop an understanding of linear programming and some of its broad applicability
5. Unfold some of the aims of the current research frontier

References

1. Kevin Leyton-Brown and Yoav Shoham, *Essentials of Game Theory*, Morgan & Claypool Publishers, 2008
2. Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V.Vazirani, *Algorithmic Game Theory*, Cambridge University Press, 2007
3. Tim Roughgarden, *Twenty Lectures on Algorithmic Game Theory*, Cambridge University Press, 2016
4. Sanjoy Dasgupta, Christos Papadimitriou, and Umesh V Vazirani, *Algorithms*, McGraw-Hill,2017

ICT 5423 APPLIED NATURAL LANGUAGE PROCESSING [4 0 0 4]

Abstract

Word Vectors, Word Window Classification, Neural Networks and Backpropagation, Dependency Parsing, Recurrent Neural Networks and Language Models, Vanishing Gradients, Fancy RNNs, Seq2Seq, Machine Translation, Attention, Transformers, Question Answering, Natural Language Generation, Conference Resolution, Large Language Models, Integrating Knowledge in Language Models

Self-Directed Learning: Vector space, Basics of neural networks

Course Outcomes

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

1. Understand the basic concepts and algorithms of natural language processing.
2. Apply neural networks for language modelling
3. Generate natural language and transform using machine learning models
4. Discover entity-emotions relationship through conflict resolution
5. Develop language models by integrating knowledge.

References

1. Daniel Jurafsky and James H.Martin, *Speech and Laguage Processing*, 3rd Edition (draft), Pearson, 2021
2. Jacob Eisenstein, *Introduction to Natural Language Processing*, The MIT Press, 2019

3. Delip Rao, and Brian McMahan, *Natural Language Processing with PyTorch*, O'Reilly Media, 2019
4. Eugene Charniak, *Introduction to Deep Learning*, The MIT Press, 2019

ICT 5412 COMPUTATIONAL SOCIAL SCIENCE [3 0 0 3]

Abstract

Introduction, Automated information extraction, Social networks, Origins and Measurement of Social Complexity, Laws of Social Complexity, Social Complexity Theories, Social Simulation Methodology, Variable-Oriented Social Simulation Models, Object-Oriented Social Simulation Models, Case study: Gender imbalance in big data, Market evolution, Natural language processing, Participatory budgeting

Self-Directed Learning: Basics of social network analysis, Familiarity with NodeXL

Course Outcomes

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

1. Use computational social science concepts to extract automated information through distinct approaches
2. Describe various laws and theories related to social complexity
3. Demonstrate the utility of social simulation models
4. Evaluate social research from the perspectives of both social science and data science
5. Create research proposals that blend ideas from social science and data science.

References

1. Claudio Cioffi-Revilla, *Introduction to Computational Social Science*, 2nd Edition, Springer, 2017
2. Tamas Rudas, and Gabor Peli, *Pathways between Social Science and Computational Social Science*, Springer, 2021
3. Frank Dignum, *Social Simulation for a Crisis*, Springer, 2021

ICT 5413 COMPUTER VISION AND APPLICATIONS [3 0 0 3]

Abstract

Introduction, Image Formation: Geometric primitives and transformations, Photometric image formation; Image Processing: Point operators, Linear filtering, Pyramids and wavelets; Model Fitting and Optimization: Scattered data interpolation, Regularization; Recognition: Instance recognition, Object detection, Semantic segmentation; Feature Detection and Matching: Points and patches, Edges and contours, Contour tracking; Image Alignment and Stitching: Pairwise alignment, Image stitching, Global alignment, Compositing; Motion Estimation: Translational alignment, Parametric motion, Optical flow, Layered motion

Self-Directed Learning: 2D transformations, 3D transformations, 3D rotations, 3D to 2D projections

Course Outcomes

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

1. Apply the concepts of image formation and image processing to solve real world problem.
2. Use appropriate model fitting and optimization techniques to solve complex engineering problem.
3. Apply suitable image recognition, stitching and feature detection techniques to solve real world problem.
4. Apply the concepts of motion estimation to solve complex engineering problem.

References

1. Richard Szeliski, *Computer Vision: Algorithms and Applications*, 2^{nd} Edition, Springer, 2021
2. David A. Forsyth and Jean Ponce, *Computer Vision: A Modern Approach*, 2nd Edition, Prentice Hall, 2011.
3. Richard Hartley, *Multiple View Geometry in Computer Vision*, 2nd Edition., Cambridge University Press, 2004

ICT 5424 DECISION INTELLIGENCE [3 0 0 3]

Abstract

Introduction to Decision Intelligence (DI), Complexity Ceiling, Buildings Decision Models, Power of Decision Model Framework, Causal and Explanatory Models in Risk Assessments, Structure of Bayesian Networks, Building and Eliciting Node Probability Tables, Numeric Variables and Continuous Distribution Functions, Decision Trees, Decision Analysis, Modeling Operational Risk, Learning from Data in Bayesian Networks

Self-Directed Learning: Decision Functions, Decision Trees, Bayesian Networks

Course Outcomes

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

1. Discuss the fundamental concepts decision intelligence
2. Implement decision intelligence models for real world applications
3. Analyze the performance of decision intelligence models
4. Demonstrate decision intelligence models for operational risk
5. Compare decision intelligence models

References

1. Lorian Pratt, *Link: How Decision Intelligence Connects Data, Actions, and Outcomes for a Better*, World,Emerald Publishing Ltd, 2019
2. Norman Fenton, and Martin Neil, *Risk Assessment and Decision Analysis with Bayesian Networks*, 2nd Editon, CRC Press, 2019.
3. L.Enrique Sucar, Eduardo F.Morales, Jesse Hoey, *Decision Theory Models for Applications in Artificial Intelligence*, IGI Global, 2011

ICT 5405 FEDERATED LEARNING [4 0 0 4]

Abstract

Introduction; Privacy-preserving Machine Learning: PPML, and secure ML; Distributed Machine Learning: Scalability Motivated DML; Horizontal Federated Learning: Architecture of HFL, Federated Averaging Algorithm; Vertical Federated Learning: Architecture of VFL; Algorithms of VFL; Federated Transfer Learning; Incentive Mechanism Design for Federated Learning; Federated Reinforcement Learning; Applications of Federated Learning

Self-Directed Learning: Basics of cloud and edge computing

Course Outcomes

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

1. Make use of concepts of federated learning for real world applications.
2. Apply federated learning for privacy preserving applications.
3. Choose suitable federated learning model for the given application
4. Design federated learning using incentive mechanisms
5. Apply federated learning model for secure reinforcement learning

References

1. Qiang Yang, Yang Liu, Yong Cheng, Yang Kang, Tianjian Chen, and Han Yu, *Federated Learning*, Morgan & Claypool Publishers, 2019
2. Qiang Yang, Lixin Fan and Han Yu, Editors, *Federated Learning: Privacy and Incentive*, Springer, 2020
3. Muhammad Habib ur Rehman, and Mohamed Medhat Gaber, Editors., *Federated Learning Systems*, Springer, 2021
4. J.Morris Chang, Di Zhuang, and G.Dumindu Samaraweera, *Privacy-Preserving Machine Learning*, Manning, 2022

ICT 5425 FULL STACK MACHINE LEARNING [3 0 0 3]

Abstract

Introduction; Organizing Machine Learning Project: Complexity Estimation of a ML Project, Defining the Goal of ML Project, Structuring a ML Team; Data Preparation: Questions about Data, Data Partitioning; Feature Engineering: Stacking Features, Properties of Good Features; Supervised Model Training; Model Evaluation; Model Deployment; Introduction to MLOps; MLOps Foundations; MLOps for Containers and Edge Devices; Continuous Delivery for Machine Learning Models; AutoML, Monitoring and Logging

Self-Directed Learning: Familiarity with DevOps, Version control systems

Course Outcomes

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

1. Demonstrate the concepts data preprocessing and feature engineering using any dataset.

2. Apply supervised training models for real-time problems.
3. Apply various evaluation models and deployment methods for real-time problems.
4. Apply the MLOps pipeline in a range of real-time applications.
5. Understand the usage of containers, continuous delivery and logging in MLOps

References

1. Andriy Burkov, *Machine Learning Engineering*, True Positive Inc., 2020
2. Noah Gift, and Alfredo Deza, *Practical MLOps: Operationalizing Machine Learning Models*, O'Reilly, 2021.
3. Emmanuel Raj, *Engineering MLOps*, Packt Publishing Ltd, 2021.
4. Villiappa Lakshmanan, Sara Robinson, and Michael Munn, Machine Learning Design Patterns, O'Reilly, 2020.
5. Alok Kumar, *Practical Full Stack Machine Learning*, BPB Publications, 2022.

ICT 5426 REINFORCEMENT LEARNING [4 0 0 4]

Abstract

Introduction to Reinforcement Learning; Muti-armed Bandits: K-armed Bandit Problem, Action-value Methods; Finite Markov Decision Processes: Agent-Environment Interface, Goals and Rewards; Dynamic Programming: Policy Evaluation, Policy Improvement; Monte Carlo Methods: Monte Carlo Estimation of Action Values, Monte Carlo Control; Temporal Difference Learning: TD Prediction, Advantages of TD Prediction Methods; On-policy Prediction with Approximation: Value-function Approximation, Prediction Objective; On-policy Control with Approximation: Episodic Semi-gradient Control, Semi-gradient n-step Sarsa; Policy Gradient Methods: Policy Approximation, Policy Gradient Theorem

Self-Directed Learning: Markov Process, Markov Chain, MDPs

Course Outcomes

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

1. Develop a reinforcement learning (RL) system that knows how to make automated decisions
2. Demonstrate how RL relates and fits into the broader umbrella of machine learning, deep learning, supervised and unsupervised learning
3. Examine the space of RL algorithms (Temporal Difference learning, Monte Carlo, Sarsa, Q-learning, Policy Gradient, Dyna, and more)
4. Formulate and implement a solution for a given RL scenario

References

1. Richard S.Sutton, and Andrew G.Barto, *Reinforcement Learning: An Introduction*, 2nd edition, The MIT Press, 2018

2. Marco Wiering, and Martijn van Otterlo, Eds., *Reinforcement Learning: State of the Art*, Springer, 2012
3. Peter Norvig, and Stuart Russel, *Artificial Intelligence: A Modern Approach*, 4th Edition, Pearson, 2021
4. Enes Bilgin, *Mastering Reinforcement Learning with Python*, Packt Publishing Ltd, 2020

ICT 5409 SEMANTIC WEB TECHNOLOGIES [3 0 0 3]

Abstract

The Semantic Web Activity of W3C: Overview of techniques and standards, XML with Document Type Definitions and Schemas; Describing Web Resource: RDF data models, syntax, semantics, schema, RDFS, RDF Data structures, Containers and collections; Querying Semantic Web: SPARQL matching patterns, filters, querying schemas; Ontology and Information Systems: Use of ontologies, types, design principles, methodologies; Ontology Languages: OWL2, OWL2 profiles; Logic for the Semantic Web: Predicate and Description Logics; Ontology Reasoning: Monotonic rules, Rule interchange format, Semantic web rules languages, RuleML; Ontology Design and Management: Types, purposes, creating ontology manually, reusing existing, mapping

Self-Directed Learning: RDF-data model, syntaxes, RDFS-adding semantics, RDF schema, RDF and RDF schema in RDFS

Course Outcomes

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

1. Demonstrate the understanding of the knowledge representation formalisms in use on the Semantic Web.
2. Build Queries using SPARQL to query the Semantic Web.
3. Represent and reason ontologies using OWL.
4. Apply ontology engineering approaches to develop ontologies.

References

1. Grigoris Antoniou, Paul Groth, Frank van Harmelen, Rinke Hoekstra, *A Semantic Web Primer*, 3rd edition, The MIT Press, 2012.
2. Peter Szeredi, Gergely Lukacsy, Tamas Benko, and Zsolt Nagy, *The Semantic Web Explained*, Cambridge University Press, 2014
3. Liyang Yu, *Introduction to the Semantic Web and Semantic Web Services*, CRC Press, 2019
4. Elisa F.Kendall, Deborah L.McGuinness, Ying Ding, and Paul Groth, *Ontology Engineering*, Morgan & Claypool Publishers, 2019

ICT 5427 TIME SERIES ANALYSIS AND FORECASTING [4 0 0 4]

Abstract
Introduction, R Language; Statistics for Forecasting: Graphical Displays, Numerical Description of Time Series Data; Regression Analysis and Forecasting: Least Square Estimation in Linear Regression Models, Statistical Inference in Linear Regression; Exponential Smoothing Methods: First and Second-Order Exponential Smoothing, Modeling Time Series Data; ARIMA Models: Linear Models for Stationary Time Series, Finite Order Moving Average Processes; Transfer Functions and Intervention Models: Transfer Function Models, Transfer Function-Noise Models

Self-Directed Learning: Basics of R language programming

Course Outcomes

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

1. Make use of concepts of time series data and represent real world problems.
2. Build regression analysis and forecast model for the given data
3. Build different time series models for forecasting
4. Apply intervention analysis using transfer function models
5. Demonstrate working of time series and forecasting algorithms using R language

References

1. Douglas C.Montgomery, Cheryl L.Jennings, and Murat Kulahci, *Introduction to Time Series Analysis and Forecasting*, 2nd Edition, Wiley-Interscience, 2015.
2. Peter J.Brockwell, and Richard A.Davis, *Introduction to Time Series and Forecasting*, 3rd Edition, Springer, 2016.
3. Jan G.De Gooijer, *Elements of Nonlinear Time Series Analysis and Forecasting*, Springer, 2018.
4. Galit Shmueli, and Kenneth C. Lichtendahl Jr, *Practical Time Series Forecasting with R*, 2nd Edition, Axelrod Schnall Publishers, 2016
5. Rob J Hyndman, and George Athanasopoulos, *Forecasting: Principles and Practice,* 3rd Edition, OTexts, 2021

ICT 5428 STATISTICAL MACHINE LEARNING [4 0 0 4]

Abstract
Statistical Machine Learning Framework, Formal Machine Learning Algorithms, Convergence of Time-Invariant Dynamical Systems, Batch Learning Algorithm Convergence, Random Vectors and Random Functions, Stochastic Sequences, Probability Models for Data Generation, Monte Carlo Markov Chain Algorithm Convergence, Adaptive Learning Algorithm Convergence, Statistical Learning Objective Function Design, Simulation Methods for Evaluating Generalization, Model Selection and Evaluation

Self-Directed Learning: Averages, Deviations, Bounds: Markov's inequality, Chebyshev's inequality; Probability, Sampling

Course Outcomes

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

1. Use statistical concepts to machine learning problems
2. Identify suitable statistical tool to analyze a machine learning algorithm
3. Experiment with deterministic and stochastic learning algorithms
4. Interpret the results of statistical analysis for a given problem
5. Statistically compare the performance of learning algorithms.

References

1. Richard Golden, *Statistical Machine Learning*, CRC Press, 2020

Jerome Friedman, Robert Tibshirani, and Trevor Hastie, *The Elements of Statistical Learning*, 2nd Edition, Springer, 2017 Masashi Sugiyama, *Introduction*

OPEN ELECTIVES

ICT 5301 APPLIED GAME THEORY [3 0 0 3]

Abstract

Introduction, Single-Person Decision Problem, Uncertainty and Time, Rationality and Common Knowledge, Nash Equilibrium, Mixed Strategies, Credibility and Sequential Rationality, Multistage Games, Repeated Games, Strategic Bargaining, Bayesian Games, Auctions and Competitive Bidding, Mechanism Design, Applications of Game Theory

Self-Directed Learning: Case studies on applications of non-cooperative, and cooperative game theory

Course Outcomes

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

1. Identify strategic situations and represent them as games
2. Solve simple games using various techniques
3. Recommend and prescribe which strategies to implement
4. Develop mechanisms to elicit the required response
5. Analyze engineering situations using game theoretic techniques

References

1. Steven Tadelis, Game Theory: An Introduction, Princeton University Press, 2013
2. Vladimir Mazalov, *Mathematical Game Theory and Applications*, Wiley, 2014
3. Hans Peters, *Game Theory: A Multi-Leveled Approach*, 2nd Edition, Springer, 2015
4. Dario Bauso, *Game Theoy with Engineering Applications*, SIAM, Philadelphia, 2016

ICT 5302 BLOCKCHAIN TECHNOLOGIES [3 0 0 3]

Abstract

Introduction to technology stack: Blockchain, protocol, understanding how blockchain works. Introduction to consensus model. Architecture of decentralized application, Dapps development process and command, application model for Dapp, introduction to Dapp development environment. Introduction to smart contracts and its development environment. Introduction to blockchain applications in different domains like government, health and genomics.

Self-Directed Learning: Cryptography concepts, Hashing techniques, Merkle tree

Course Outcomes

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

1. Apply concepts of Blockchain technologies to build application.
2. Analyze different components of blockchain ecosystem, the elements of trust in a blockchain: validation, verification, and consensus
3. Examine the relevance of Blockchain technology in different applications/ usecases
4. Develop smart contracts in Ethereum framework

References

1. Melani Swan, Blockchain: Blueprint for a New Economy (1e), O'Reilly Media, 2015.
2. Paul Vigma, Michael J. Casey, The Truth Machine: The Blockchain and the Future of Everything (1e), St Martin's Press, 2018.
3. Daniel Drescher, Blockchain Basics: A Non-Technical Introduction in 25 Steps (1e), Apress, 2017.
4. Elad Elrom, The Blockchain Developer: A Practical Guide for Designing, Implementing, Publishing, Testing, and Securing Distributed Blockchain-based Projects (1e) , Apress ,2019
5. Imran Bashir, Mastering Blockchain: Distributed Ledger Technology, decentralization, and smart contracts explained(2e), Packt Publishing Ltd, 2018.
6. Bellaj Badr, Richard Horrocks, Xun (Brian) Wu, Blockchain By Example: A developer's guide to creating decentralized applications using Bitcoin, Ethereum, and Hyperledger, Packt Publishing Limited, 2018.

ICT 5303 CYBER SECURITY AND CYBER LAW [3 0 0 3]

Abstract

Introduction to Information, Network and System Security, Encryption techniques, Message Integrity and Message Authentication, Digital Signature, Key Management, User Authentication. Web security model: Browser security model including same-origin policy, Client-server trust boundaries, Session management, authentication: Single sign-on, HTTPS and certificates. Application vulnerabilities and defenses: SQL injection, XSS, CSRF. Client-side security: Cookies security policy, HTTP security extensions, Plugins, extensions, and web apps, Web user tracking, Server-side security tools, e.g. Web Application Firewalls (WAFs) and fuzzers. Cybercrime, Cybercrime investigation, Laws and ethics

Self-Directed Learning: SQL injection, XSS, CSRF, Web Application Firewalls

Course Outcomes
By the end of this course, the student should be able to

1. Understand the symmetric and asymmetric cryptographic algorithms.
2. Describe common types of vulnerabilities and attacks in web applications, and defenses against them.
3. Understand client side and server side security concepts and tools
4. Propose and design security algorithm for a particular application
5. Understand cybercrimes, cybercrime investigation, Laws and ethics.

References:

1. Mayank Bhushan, *Fundamentals of cybersecurity*, BPB publications, 2017
2. Raef Meeuwisse, *Cyber Security for Beginners*, 2015
3. Nilakshi Jain and Ramesh Menon, *Cyber Security and Cyber Laws*, Wiley, 2020
4. Rolf Oppliger, *Security Technologies for the World Wide Web* (2e), Artech House, 2002.
5. Seth Fogie, Jeremiah Grossman, Robert Hansen and Anton Rager, *XSS Attacks: Cross Site Scripting Exploits and Defense*, Syngress, 2007.
6. Justin Clarke et.al.,*SQL Injection Attacks and Defense (2e)*, Syngress, 2012.

7. DafyddStuttard, and Marcus Pinto, *The Web Application Hacker's Handbook: Finding and Exploiting Security Flaws* (2e), Wiley, 2011.

ICT 5304 REAL TIME SYSTEMS [3 0 0 3]

Abstract:

Introduction to Real Time Systems, Resource management, Commonly used approaches for real time scheduling-static scheduling, priority driven scheduling, RM and DM algorithms, Aperiodic jobs and scheduling, Computation of average response time, Various servers: Deferrable, Sporadic etc. Bandwidth computation, Resource access protocols: various resources access protocols and features, Advantages and drawbacks, Priority ceiling protocols and its use in dynamic priority systems, multiprocessor scheduling, Task assignment and conditions, Faults and fault handling, Redundancy and handling redundancy, Real time communication

Self-Directed Learning: Motivation and features of real-time operating systems

Course Outcomes

By the end of this course, students should be able to:

1. Understand basic real time system model
2. Comprehend the salient features of various RTOS
3. Analyse various real time scheduling algorithms
4. Apply the principles of resource access protocols

References:

1. Jane W.S.Liu, *Real Time Systems*, Pearson Edition-2006.
2. C.M Krishna and K.G Shin, *Real Time Systems*, 1st Edition, McGraw Hill Education, 2017
3. Philip A Laplante and Seppo J Ovasaka , *Real-Time Systems Design and Analysis; Tools for the Practitioners* , 4th Edition, Wiley, 2013