



IISc's Knowledge and E-Learning Network

(IKEN)

M.Tech (Online) Sponsored Degree Programme

Scheme of Instructions

Artificial Intelligence (AI)

The Master of Technology (Online) programme in Artificial Intelligence is offered by the Division of Electrical, Electronics, and Computer Sciences (EECS). The vision of this programme is to impart rigorous training in the foundations and deep technology of Artificial Intelligence to early-career professionals with 2-8 years of experience to upskill them to become technology and business leaders in information-driven enterprises. These learnings are coupled with a unique capstone project that applies the learnings to a hands-on project relevant to the industry. Faculty members from the Division of EECS, the Department of CDS, and the RBCCPS will offer contemporary courses in AI through online lectures and tutorials, and will provide mentorship on capstone projects.

Programme Structure

- **Core Courses (16 Credits):** These are typically taken in the first and second semesters.
 - E0 270o: Machine Learning (3:1)
 - E1 220o: Linear Algebra (3:1)
 - E1 252o: Linear and Non-Linear Optimization (3:1)
 - E2 202o: Random Processes (3:1)

- **Elective Courses (at least 21 Credits):**
 - CP 214o: Foundations of Robotics (3:1)
 - DA 201o: Applied AI: Building Practical and Scalable ML Systems (3:1)
 - DA 204o: Data Science in Practice (3:1)
 - DA 219o: Quantum Computing Methods: Theory and Applications (3:1)
 - DA 245o: Linear Optimization and Network Science (3:1)
 - DS 216o: Applied Artificial Intelligence in Healthcare (3:1)
 - DS 265o: Deep Learning for Visual Analytics (3:1)
 - DS 285o: Tensor Computations for Data Science (3:1)
 - E0 251o: Data Structures and Graph Analytics (3:1)
 - E0 259o: Data Analytics (3:1)
 - E0 295o: Introduction to Cryptography (3:1)
 - E1 277o: Reinforcement Learning (3:1)
 - E1 285o: Advanced Deep Representation Learning (3:1)
 - E1 286o: Deep Generative Models (3:1)
 - E1 294o: Edge and Cloud Systems for Machine Learning Algorithms (3:1)
 - E1 316o: Deep Learning for Robotics (2:1)
 - E2 276o: Process Technology and System Engineering (3:1)
 - E9 241o: Digital Image Processing (3:1)
 - E9 246o: Advanced Topics in Image Processing (3:1)

Students may also take courses from any of the three streams as an elective. These are the minimum number of elective credits. More may be taken as well.

- **Project (27 Credits):**

Students can start this three-term project for 27 credits after successfully completing all the core courses. Students will complete 3 project credits in the 1st term to identify the topic, 12 credits in the 2nd term followed by a mid-term evaluation, and the remaining 12 project credits the 3rd term for the final-term evaluation. Student will propose the topic in consultation with their guide from within the organization, and an IISc faculty mentor will approve project goals. The faculty mentor will offer high-level feedback on the project and its progress, and coordinate the evaluations, while the in-house company guide will offer active feedback and close support. The evaluation committee, which includes the faculty mentor and company guide, is appointed by the Programme Curriculum Committee (PCC).

Data Science and Business Analytics (DSBA)

The Master of Technology (Online) programme in Data Science and Business Analytics is offered by the Division of Interdisciplinary Sciences. It is designed for early-career professionals with 2-8 years of experience to upskill them to become technology and business leaders in information-driven enterprises. The designed coursework establishes the foundations of data science, trains on data engineering and machine learning techniques, and imparts practical business analysis skills. These learnings are coupled with a unique capstone project that applies the learnings to a hands-on project relevant to the industry. Faculty members from the Departments of CDS, Management Studies, CiSTUP, and RBCCPS lead the courses through online lectures and tutorials.

Programme Structure

- **Core Courses (12 Credits):** These are typically taken in the first and second semesters.
 - DA 201o: Applied AI: Building Practical and Scalable ML Systems (3:1)
 - DA 204o: Data Science in Practice (3:1)
 - DA 231o: Data Engineering at Scale (3:1)
- **Elective Courses (at least 20 Credits):**
 - DA 202o: Introduction to Data Science (3:1)
 - DA 203o: Introduction to Computing for AI & Machine Learning (3:1)
 - DA 212o: MLOps at Scale (3:1)
 - DA 218o: Probabilistic Machine Learning: Theory and Applications (3:1)
 - DA 219o: Quantum Computing Methods: Theory and Applications (3:1)
 - DA 224o: Practical Machine Learning (3:1)
 - DA 225o: Deep Learning (3:1)
 - DA 226o: Financial Analytics (3:1)
 - DA 227o: Data Mining (3:1)
 - DA 245o: Linear Optimization and Network Science (3:1)
 - DS 216o: Applied Artificial Intelligence in Healthcare (3:1)
 - DS 261o: Artificial Intelligence for Medical Image Analysis (3:1)
 - DS 285o: Tensor Computations for Data Science (3:1)
 - E0 295o: Introduction to Cryptography (3:1)
 - E1 220o: Linear Algebra (3:1)
 - E1 252o: Linear and Non-Linear Optimization (3:1)
 - E1 285o: Advanced Deep Representation Learning (3:1)
 - E1 286o: Deep Generative Models (3:1)
 - E1 294o: Edge and Cloud Systems for Machine Learning Algorithms (3:1)
 - E1 316o: Deep Learning for Robotics (2:1)
 - E2 276o: Process Technology and System Engineering (3:1)
 - E9 246o: Advanced Topics in Image Processing (3:1)

Students may also take courses from any of the three streams as an elective. These are the minimum number of elective credits. More may be taken as well.

- **Project (32 Credits):**

Students can start this three-term project for 32 credits after successfully completing all the core courses. Students will complete 20 project credits over two terms followed by a mid-term evaluation, and the remaining 12 project credits in an exclusive semester for the final-term evaluation. Student will propose the topic in consultation with their guide from within the organization, and an IISc faculty mentor will approve project goals. The faculty mentor will offer high-level feedback on the project and its progress, and coordinate the evaluations, while the in-house company guide will offer active feedback and close support. The evaluation committee, which includes the faculty mentor and company guide, is appointed by the Programme Curriculum Committee (PCC).

Electronics and Communication Engineering (ECE)

The Master of Technology (Online) programme in Electronics and Communication Engineering is offered by the Division of Electrical, Electronics, and Computer Sciences (EECS). The programme is designed for early-career professionals with 2-10 years of experience to strengthen their fundamentals and expose them to cutting-edge topics in the fast-changing areas of communications, networks, signal processing and information sciences, and high-frequency circuits and systems. The coursework contains core courses, which provide the necessary foundation, and several elective courses, which provide exposure to the state-of-the-art and advanced material. These learnings are coupled with a semester-long capstone project that enables the student to apply the knowledge gained to a project relevant to the industry. The courses are taught online by faculty members from the Department of Electrical Communication Engineering (ECE).

Programme Structure

- **Core Courses (16 Credits):** These are typically taken in the first and second semesters.
 - E1 220o: Linear Algebra (3:1)
 - E1 245o: Statistical Inference for Engineers and Data Scientists (3:1)
 - E2 202o: Random Processes (3:1)

- **Elective Courses (20 Credits):**
 - DA 201o: Applied AI: Building Practical and Scalable ML Systems (3:1)
 - DA 204o: Data Science in Practice (3:1)
 - DA 219o: Quantum Computing Methods: Theory and Applications (3:1)
 - DA 245o: Linear Optimization and Network Science (3:1)
 - DS 216o: Applied Artificial Intelligence in Healthcare (3:1)
 - DS 285o: Tensor Computations for Data Science (3:1)
 - E0 295o: Introduction to Cryptography (3:1)
 - E1 252o: Linear and Non-Linear Optimization (3:1)
 - E1 285o: Advanced Deep Representation Learning (3:1)
 - E1 286o: Deep Generative Models (3:1)
 - E1 294o: Edge and Cloud Systems for Machine Learning Algorithms (3:1)
 - E1 316o: Deep Learning for Robotics (2:1)
 - E2 203o: Wireless Communications (3:1)
 - E2 251o: Communication Systems Design (3:1)
 - E2 276o: Process Technology and System Engineering (3:1)
 - E2 287o: Communication Networking Lab (1:1)
 - E3 280o: Semiconductor Devices for Nanoelectronics (3:1)
 - E8 204o: Antenna Theory and Practice (3:1)
 - E8 242o: Radio Frequency Integrated Circuits and Systems (3:1)
 - E9 246o: Advanced Topics in Image Processing (3:1)

Students may also take courses from any of the three streams as an elective. These are the minimum number of elective credits. More may be taken as well.

- **Project (28 credits):**

This involves a two-term project, with 10 credits in the first term followed by a mid-term evaluation, and 18 credits in the second-term with a final-term evaluation. The second term project is taken up only after all courses are completed in prior semesters. Student will propose the topic in consultation with their guide from within the organization, and an IISc faculty mentor will approve project goals. The faculty mentor will offer high-level feedback on the project and its progress, and coordinate the evaluations, while the in-house company guide will offer active feedback and close support. The evaluation committee, which includes the faculty mentor and company guide, is appointed by the Programme Curriculum Committee (PCC).

ARTIFICIAL INTELLIGENCE
(AI)

AI Core Courses

E0 270o Machine Learning (3:1)

Course Instructor

Prof. Ambedkar Dukkipati, CSA

Course Description

The aim of the course is to provide a unified view of Artificial Intelligence (AI) and Machine Learning (ML) methods. This course stresses on foundations and covers important topics in three elements of Machine Learning: supervised learning, unsupervised learning, and reinforcement learning. Further, this course provides basics of deep learning and emphasizes how deep neural networks play key role in various problems as the powerful function approximators. Assignments will involve programming and require knowledge of python.

Webpage will be made available here.

Syllabus

- *Supervised Learning*: Classification with Bayes rule, learning as optimization, linear regression, logistic regression, probabilistic view: ML and MAP estimates, gradient descent, hyperplane-based classifiers, perceptron, kernel methods, feedforward neural networks, backpropagation algorithm, CNN (Convolutional Neural Network), RNN and LSTMs
- *Unsupervised Learning*: Principal component analysis, clustering methods, undirected graphical models, MCMC (Markov Chain Monte Carlo) and Gibbs sampling, latent variable and mixture models, EM (Expectation-Maximisation) algorithm, Deep generative models
- *Reinforcement Learning*: Introduction to sequential decision making and online learning, Markov decision processes and dynamic programming

Textbooks / References

1. Pattern Recognition and Machine Learning by C. M. Bishop, 2006
2. The Elements of Statistical Learning: Data Mining, Inference and Prediction by Hastite T, Tibshirani R and Friedman J, 2009
3. Neural Networks and Learning Systems by Haykin. S, 2009
4. Deep Learning by Goodfellow, Bengio, Courville, 2017

Prerequisites: There are no formal prerequisites.

Grading

- Homework: 40%
- Mid-Term Exam: 30%
- End-Term Exam: 30%

E1 220o Linear Algebra (3:1)

Course Instructor

Prof. Sundeep Prabhakar Chepuri, ECE

Course Description

In this course, we will study the basics of linear algebra and matrix theory, with applications to engineering. The focus will be two-fold: on the beautiful mathematical theory of matrices, and their use in solving engineering problems.

Syllabus

Fundamental ideas - vector spaces, matrices, determinant, rank, etc.; Norms, error analysis in linear systems; Eigenvalues and eigenvectors; Canonical, Symmetric and Hermitian forms, matrix factorizations; Least-squares problems, generalized inverses; Miscellaneous topics/applications

Textbooks / References

1. Matrix Analysis by Horn and Johnson, Cambridge University Press
2. Matrix Theory by David Lewis, Allied Publishers
3. Matrix Computations by Golub and Van Loan, 3rd Ed., John Hopkins University Press
4. Linear Algebra and its Applications by Gilbert Strang, 3rd Ed., Harcourt Brace Janovich Pubs. (See also: <http://ocw.mit.edu/courses/mathematics/18-06-linear-algebra-spring-2010/video-lectures/>)

Prerequisites: None

Grading

- Homework Assignments (four): 40%
- Mid-Term Exam: 30%
- End-Term Exam: 30%

E1 252o Linear and Non-Linear Optimization (3:1)

Course Instructor

Prof. Chandramani Singh, ESE

Course Description

In this course, we will study the basics of linear and nonlinear optimization. We will also see several usages of optimization techniques in supervised and unsupervised learning.

Syllabus

Optimization examples, The basics - global vs local optimality; Convex sets, Convex and concave functions; First-order and second-order optimality conditions; Gradient descent methods, Conjugate gradient method, Newton method, Gradient projection method; Constrained optimization with equality and inequality constraints, Duality; Linear programming, simplex method, duality; Barrier and penalty function methods; Sub gradient descent methods; Proximal gradient descent; Augmented Lagrangian methods

Textbooks / References

1. Nonlinear Programming by D. Bertsekas, Athena Scientific, 2016
2. Linear and Nonlinear Programming by D. Luenberger and Y. Ye, Springer, 2008
3. Convex Optimization by S. Boyd and L. Vandenberghe, Cambridge University Press, 2004

Prerequisites: None

Grading

- Homework (assigned approximately once in two weeks): 40%
- Mid-Term Exam: 30%
- End-Term Exam: 30%

Homeworks will be assigned approximately once in two weeks, and a scanned copy of the solutions need to be turned in.

E2 2020 Random Processes (3:1)

Course Instructor

Prof. Aditya Gopalan, ECE

Course Description

This course is a graduate-level course on probability and stochastic processes. It is assumed that the students are familiar with multivariable calculus (functions of several variables, partial derivatives, integration in n-dimensional real Euclidean spaces) and that they have some idea of elementary probability (e.g., as part of Foundations for Business Analytics of an undergraduate course on mathematics). Some familiarity with vector spaces and matrices would be assumed. The course would be useful for first year Masters or Ph.D. students and would equip them with basic background in probability which is required for more advanced courses such as Machine Learning, Adaptive Signal Processing etc. The course is a mathematics course, and the students are encouraged to solve many problems. There would be some tutorial classes to help students with problem solving.

Syllabus

The axioms of probability theory, probability spaces, conditional probability, independence, random variables and distribution functions, continuous and discrete random variables, multiple random variables and joint distributions, conditional distributions, functions of random variables and random vectors, expectation and moments, conditional expectation, some moment inequalities, sequences of random variables and convergence concepts, laws of large numbers, sums of independent random variables and the central limit theorem, stochastic processes, stationarity and ergodicity, discrete time Markov chains, Poisson process, continuous time Markov chains, Brownian motion

Textbooks / References

1. An Introduction to Probability and Statistics by V. K. Rohatgi and A. K. M. E. Saleh
2. Introduction to Probability Theory by P. G. Hoel, S. C. Port and C. J. Stone
3. Introduction to Stochastic Processes by P. G. Hoel, S. C. Port and C. J. Stone
4. Introduction to Probability Models by S. M. Ross

Prerequisites: There are no formal prerequisites. However, students should be familiar with multivariate calculus.

Grading

- Homework Assignments (four): 40%
- Mid-Term Exam: 30%
- End-Term Exam: 30%

AI Elective Courses

CP 214o Foundations of Robotics (3:1)

Course Instructor

Prof. Shishir YNK, RBCCPS

Course Description

This graduate course will serve as an introductory robotics course for students with little/no background in mechanical systems. It will first build the necessary mathematical framework in which to understand topics such as centre of gravity and moment of inertia, friction, statics of rigid bodies, principle of virtual work, kinematics of particles and rigid bodies, impacts, Newtonian and Lagrangian mechanics, rigid body transformations, forward and inverse kinematics, forward and inverse dynamics, state space representations. Towards the end of the course, advanced topics such as rigid body collisions, and hybrid dynamical systems will also be covered.

Syllabus

Kinematics of particles and rigid bodies, statics and dynamics of rigid bodies, moment of inertia, principle of virtual work, conservation of energy and momentum, collisions, configuration space, task space, rotation groups, rigid transformations, forward and inverse kinematics, forward and inverse dynamics, holonomic and nonholonomic constraints, hybrid systems, hybrid modeling

Week 1: **Part I:** Introduction, Overview of course, overview of mechanical systems

Week 2: Free-body diagrams, constraints, friction, centre of gravity and moment of inertia

Week 3: Virtual displacement, principle of virtual work, potential energy, and equilibrium

Week 4: **Part II:** Kinematics and Dynamics of Rigid Bodies, Types of motion, force, acceleration

Week 5: Work and energy, impulse and momentum, impact

Week 6,7: Equations of motion **Part III:** Kinematics and Dynamics of Robots

Week 8,9: Configuration space, task space, rigid body transformations

Week 10,11: Manipulator kinematics, forward and inverse kinematics

Week 12: Forward and inverse dynamics

Week 13: Constrained motion, holonomic and nonholonomic systems

Textbooks / References

1. Introduction to Statics and Dynamics by Ruina, Andy and Pratap, Rudra, Oxford University Press, 2011
2. A Mathematical Introduction to Robot Manipulation by Murray, Li and Sastry, CRC Press, 1994
3. Robotics: Fundamental Concepts and Analysis by A. Ghosal, Oxford, 2006

Prerequisites: None. Basic concepts in linear algebra and programming will help.

Grading

- Homework: 50%
- Mid-Term Exam: 20%
- End-Term Exam: 30%

DA 201o Applied AI: Building Practical and Scalable ML Systems (3:1)

Course Instructor

Prof. Sashikumar Ganesan, CDS

Course Description

This four-credit course will be offered as a core course for the M.Tech. (Online) DSBA program in place of DA 203o Introduction to Computing for AI & Machine Learning 3:1. The course is designed as a graduate-level (200-series) course.

This course offers a comprehensive introduction to computational thinking with a focus on its practical applications in Artificial Intelligence and Machine Learning (AI & ML). Students will gain a thorough understanding of machine learning techniques and develop skills to design ML systems suitable for production-ready applications.

Syllabus

Python Programming, Parallel Computing, and Machine Learning Tools: Python Environment: Introduction to Python programming, setting up the Python environment, Using Shell commands, Jupyter Notebook, Colab, and VS Code, Python Installation, pip, and virtual environments. Parallel Computing Fundamentals: Understanding computer performance, Caches and Cache optimization techniques, Roofline analysis, Types of parallelism, Shared parallelism, Distributed parallelism, GPUs in Machine Learning, Introduction to GPUs and their architecture, CUDA programming basics. Machine Learning Tools: Pandas, NumPy, Matplotlib, and Seaborn, Understanding shell scripting, Piping commands, awk, sed, and Regular Expressions, Data imputations, Data cleanup with Pandas, Data cleanup with Excel.

Applied Machine Learning: Supervised Learning: Linear Regression, Polynomial Regression, Underfitting and Overfitting, Regularization Methods, Regression Metrics. Classification: Logistic Regression, Loss Functions for Classification, Confusion Matrix, Classification Metrics. Ensemble Classification: Random Forest, Decision Trees, Gini Index, Bagging and Pruning in Random Forest, Boosting, Gradient Boosting, XGBoost. Unsupervised Learning: K-Means Clustering and Elbow Method, Clustering Metrics, Feature Selection and Feature Extraction, Principal Component Analysis (PCA), Singular Value Decomposition (SVD), Linear Discriminant Analysis (LDA), Explained Variance (PCA). Neural Networks: Types of Neural Networks, Single/Multi-layered Perceptron, Activation and Loss Functions, Convolutional Neural Networks, Introduction to Deep Learning.

MLOps: Introduction to MLOps, Overview of MLOps lifecycle. Cloud Infrastructure for MLOps: Cloud services overview, Cloud account setup, Cloud data storage, Compute instance setup. Docker for MLOps: Introduction to Docker, Containerization, Creating Dockerfile for ML model, Building and running Docker container locally, Pushing Docker image to a container registry. ML Model Training and Deployment: Model training and deployment, Model scaling and automation. MLOps Pipelines: Overview of MLOps pipelines, Building an ML pipeline, Adding data validation and error handling to pipeline, Monitoring, and visualizing pipeline. Model Monitoring and Management: Importance of model monitoring and management, Creating model evaluation metrics, Automated model retraining and redeployment. Best Practices for MLOps: Best practices for managing ML models, Security and compliance considerations, Cost optimization strategies, Future trends.

Textbooks / References

1. Learn Python 3 the hard way: A very simple introduction to the terrifyingly beautiful world of computers and code by Shaw, Zed A, Addison-Wesley Professional, 2017
2. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow by Aurélien Géron, 3rd Edition, O'Reilly Media, Inc. 2022
3. Practical MLOps: Operationalizing Machine Learning Models by Noah Gift & Alfredo Deza, Shroff/O'Reilly; First edition, 2021

Prerequisites: Basic knowledge of programming

Grading

- Homework: 30%
- Mid-Term Exam: 30%
- Final Project: 20%
- End-Term Exam: 20%

DA 204o Data Science in Practice (3:1)

Course Instructor

Prof. Pandarasamy Arjunan, RBCCPS

Course Description

This is a core course for the DSBA stream and an elective for non-DSBA streams. This course provides an introduction to using Data Science, including Machine Learning and Artificial Intelligence, in practice. At the end of the course, students will be confident in identifying and solving end-to-end data science problems in practical applications.

Syllabus

Data Science Fundamentals: Identifying and framing a data science problem in different fields; Data - Types, Pre-processing; Different types of Analytics; Introduction to Machine Learning, Artificial Intelligence; Is ML/AI the right tool for your problem; Stakeholder Discussion Guidelines; End-to-end Problem Solving through a 6-Step Data Science Process.

Exploratory Data Analysis: Math Foundations of Probability and Statistics, Hypothesis Testing, How much data is sufficient data, Data Distributions, Imputation, Outlier handling

Data Science for Tabular Data: CART Algorithm, Random Forest Models, Gradient Boosted Models (XGBoost, CatBoost, LightGBM), Feature Importance and Selection, Development-Testing Paradigm, Cross Validation

Deep Learning for Computer Vision and Natural Language Processing: From Linear Regression to Neural Networks from fundamentals, Matrix and Tensor algebra for Neural Networks, Basics of Stochastic Gradient Descent and Backpropagation, Hyperparameter Tuning, Different types of Layers, NN as data-processing pipelines, Practical computer vision with Transfer Learning, Natural Language Processing with Bag of Words models and sequence transformers

Programming for Data Science, ML and AI: Python (NumPy, Pandas), Scikit Learn, TensorFlow, Keras

Textbooks / References

1. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow by Aurélien Géron, 3rd Edition, O'Reilly Media, Inc., 2022
2. A Hands-On Introduction to Data Science by Shah, Chirag, Cambridge University Press, 2020
3. Learn python 3 the hard way: A very simple introduction to the terrifyingly beautiful world of computers and code by Shaw, Zed A, Addison-Wesley Professional, 2017
4. Mathematics for Machine Learning by Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong, Cambridge University Press, 2020 (<https://mml-book.github.io>)

Prerequisites: Basic knowledge of mathematics

Grading

- Programming Assignments: 20%
- Course Project: 20%
- Biweekly Quiz: 30%
- End-Term Exam: 30%

DA 219o Quantum Computing Methods: Theory and Applications (3:1)

Course Instructor

Prof. Phani Motamarri, CDS

Course Objectives

Quantum computing (QC) has an immense potential to drastically revolutionize industries such as finance, healthcare, AI and automotive over the next several years. This course will give a basic introduction to quantum computing by exposing the participants to the underlying key ideas at the intersection of quantum mechanics, linear algebra, and computing. Additionally, the course will introduce the inner workings of popular quantum algorithms, including how to build, manipulate quantum gates, and circuits using quantum computing toolboxes.

Syllabus

Introduction: Need for quantum computing, quantum vs classical mechanics, quantum vs classical computing, quantum supremacy vs quantum advantage, types of quantum hardware, sneak-peak into a few industrial case studies.

Mathematical Foundations of Quantum Computing: Dirac notation for vectors (Kets, Bra), Linear independence of kets(vectors), basis kets(vectors), linear vector space, orthonormal vectors, inner products, matrices viewed as linear transformations, matrix-vector products, matrix-matrix multiplication, unitary and Hermitian matrices, operators in the quantum world (Hadamard, Pauli X, Y, Z operators), the action of operator on a vector, outer products of vectors, tensor products, the inverse of a matrix, eigenvalue problems, the relevance of these concepts in quantum computing.

Fundamentals of Quantum Computing: Quantum Bit (Qubit), Quantum superposition (qubit state-space), conservation of probabilities, basis transformations, Observables, Time evolution of the quantum system, quantum measurement (eg: spin measurement -via- Stern-Gerlach experiment), single-qubit gates, quantum circuit, n-qubit state space, quantum gates, quantum entanglement, Bell states (EPR states), Bell measurement.

Development Libraries for Quantum Computer Programming: Quantum Computing simulators, Introduction to Qiskit (Quantum Development Kit), Hands-on tutorials for implementation of Quantum Gates, Circuits.

Quantum Algorithms: No-cloning theorem, Quantum algorithm construction, Quantum teleportation, Entanglement swapping, Complexity of algorithms, Deutsch-Jozsa algorithm (1-qubit input, n-qubit input cases), Bernstein-Vazirani algorithm, Grover's algorithm, Introduction to Quantum machine learning, Hands-on sessions for implementation of Quantum algorithms in Qiskit

Textbooks / References

1. Quantum Computing for Everyone by Chris Bernhardt, MIT Press, Cambridge, Massachusetts, 2019
2. Quantum Computing: An Applied Approach by Jack D. Hidary, Springer, 2021, (Second Edition)
3. Quantum Computation and Quantum Information by Michael A Nielsen, Isaac L. Chuang, Cambridge University Press, 2010

Prerequisites: Basic linear algebra and probability; Some proficiency in Python

Grading

- Homeworks and Projects through Qiskit: 50%
- Mid-Term Exam: 20%
- End-Term Exam: 30%

DA 245o Linear Optimization and Network Science (3:1)

Course Instructor

Prof. Tarun Rambha, Civil Engineering

Course Description

This course is aimed at:

- (1) formulating practical problems as optimization models and using efficient algorithms for solving them, and
- (2) understanding how dynamic cyber and physical networks evolve.

We will particularly focus on linear and network optimization models which may have continuous or integer variables. These models have many applications in areas such as mobility, scheduling, energy, manufacturing, e-commerce, and logistics. A few hands-on sessions will also introduce you to solvers and visualization tools relevant to the course such as CPLEX/Gurobi, NetworkX, Gephi, and OR tools.

Syllabus

Introduction to Linear Programming (LP); Geometry of LPs; Simplex Method; Duality; Large scale optimization and applications (Column generation, Dantzig Wolfe decomposition, Benders decomposition); Introduction to Networks; Shortest paths (Label setting and label correcting methods, A* algorithm, Contraction hierarchies); Max flows and Min cost problems (Augmenting path method, Cycle cancelling and successive shortest path methods); Integer programs (Branch and bound and cutting plane method); Traveling salesman and Vehicle routing problems; Random networks and centrality (Small worlds, power laws, scale-free properties); Evolution of networks (Preferential attachment); Spreading phenomenon (Epidemics and contact networks).

Textbooks / References

1. Introduction to Linear Optimization (Vol. 6, pp. 479-530) by D. Bertsimas & J.N. Tsitsiklis (1997), Belmont, MA: Athena Scientific
2. Linear Programming and Network Flows by M.S. Bazaraa, J.J. Jarvis & H.D. Sherali (2011), John Wiley & Sons
3. Network Flows: Theory, Algorithms, and Applications by R.K. Ahuja, T.L. Magnanti & J.B. Orlin (1993), Pearson
4. Network Science by A.L. Barabási (2016), Cambridge University Press

Grading

- 5 Quizzes: 30%
- Project: 30%
- End-Term Exam: 40%

DS 216o Applied Artificial Intelligence in Healthcare (3:1)

Course Instructor

Prof. Vaanathi Sundaresan, CDS

Learning Objectives

1. To understand the machine learning (ML) concepts including deep learning and choose appropriate techniques/methods for various tasks aiming towards various clinical/healthcare applications.
2. To be able to build a data analytics pipeline suited for various real-world applications for various modalities – e.g., audio/speech/sensory signal processing, text data, image analysis applications such as object detection, tracking or counting, etc.
3. Evaluate the performance of the method with respect to a gold standard target and analyze the competency of the method.
4. Statistical analysis of results (for better analysis of data) for various tasks for given the population/dataset size in the real-world scenarios. This would provide a comprehensive skillset required for tool development, testing and deployment in healthcare industry.

Syllabus

- **ML Concepts:** Introduction to data and types of learning, Expectation maximization methods, clustering; Representation learning; ML classifiers - kernel-based methods, ensemble methods: decision trees, Bayesian networks - hidden Markov models, Conditional random fields, dimensionality reduction, Deep learning – introduction to feedforward networks, feature saliency and visualization, convolutional neural networks, encoder-decoder models, graph-based models, generative models.
- **ML Applications in Healthcare:** Data availability and regulatory framework for AI in healthcare, Template matching, correlation – audio/speech signals; Regression and classification on publicly available speech/sensory/biological signal data, image segmentation, classification & disease prognosis - machine learning classifiers, feature-based and rule-based decision making, image and text correlation, uncertainty estimation, semi-/self-supervised learning.
- **Evaluation of Analysis Tasks:** Evaluation metrics, segmentation evaluation metrics (IoU, Dice, Jaccard indices, Hausdorff distance measures), classification evaluation metrics (confusion matrix, sensitivity, specificity, accuracy), registration metrics (MSE, MAE).
- **Statistical Evaluation:** Testing statistical significance of ML applications: Review of hypothesis testing basic, permutation tests, effect of sample size, statistical power, parametric and non-parametric tests, Shapiro-wilks test, t-tests. Statistical evaluation of ML applications: Descriptive statistics (mean, standard deviation, median, confidence interval, IQR), Unpaired and paired t-tests, one-way and repeated measures ANOVA.

Textbooks / References

1. Pattern Recognition and Machine Learning by C.M. Bishop, Springer, 2006
2. Deep Learning by I. Goodfellow, Y. Bengio and A. Courville, 2016
3. The Elements of Statistical Learning by Jerome H. Friedman, Robert Tibshirani and Trevor Hastie, Springer, 2001
4. Design of Experiments: Statistical Principles of Research Design and Analysis by R.O. Kuehl, 2000
5. Research papers, material/notes provided by instructor

Prerequisites: Basic knowledge in linear algebra, probability; Proficiency in Python coding

Grading

- Assignments: 15% (best 1 out of 2)
- Mini ML Application Project: 30%
- Mid-Term Exam: 25%,
- End-Term Exam: 30%

DS 265o Deep Learning for Visual Analytics (3:1)

Course Instructor

Prof. Venkatesh Babu, CDS

Course Description

In the recent years, Deep Learning has pushed to boundaries of research in many fields. This course focuses on the application of Deep Learning in the field of Visual Analytics. It starts with the basics of Deep Learning, which are built on top of various concepts from Image Processing and Machine Learning. The second part of the course deals with the various flavours of Deep Learning in Computer Vision, such as Generative Models, Recurrent Models and Adversarial Robustness etc.

Syllabus

Basics of machine learning and computer vision, CNN basics, Loss function and back propagation, Object Recognition, Detection and Segmentation, Recurrent Neural Networks, LSTM, Generative Adversarial Networks (GANs), Self-supervised learning, Transformers, Explainable AI, Adversarial Robustness of Deep models

Textbooks / References

1. Dive into Deep learning by Aston Zhang, Zachary C. Lipton, Mu Li, and Alexander J. Smola (Online)
2. Recent research papers

Prerequisites: Basic knowledge of machine learning and image processing

Grading

- Projects: 35%
- Assignments: 25%
- Mid-Term Exam: 20%
- End-Term Exam: 20%

DS 285o Tensor Computations for Data Science (3:1)

Course Instructor

Prof. Ratikanta Behera, CDS

Course Description

This course is an introduction to tensor computations, focusing on theory, algorithms, and applications of tensor decompositions to data sciences. In the era of Big Data, Artificial Intelligence, and Machine Learning, we are faced with the need to process multiway (tensor-shaped) data. These data are mostly in the three or higher order dimensions, whose order of magnitude can reach billions. Huge volumes of multi-dimensional data are a great challenge for processing and analyzing; the matrix representation of data analysis is not enough to represent all the information content of the multiway data in different fields. Further, the importance of being able to decompose a tensor is (at least) two-fold. First, finding the decomposition provides hidden information about the data at hand, and second, having a concise decomposition of the tensor allows us to store it much more efficiently.

The course will provide an understanding of tensor operations and decomposition with various applications, including image deblurring, image compression, neural network, and solving high dimensional partial differential equations.

Learning Outcomes

- Students will learn basic understanding of the theoretical foundations of tensors computation.
- Students will choose efficient tensor decomposition for solving a specific problem.
- Students will learn efficient algorithms for tensor operations; including multiplication, decomposition, and inverse of tensors.
- Students will learn how to implement and use tensor computation in data sciences; including image deblurring, image compression, and solving high dimensional partial differential equations.
- Students will learn the basic tensor computation in the neural networks.

Syllabus

Unit-1 Fundamentals: Basic concepts of matrix properties: norms, rank, trace, inner products, Kronecker product, similarity matrix. Fast Fourier transform, diagonalization of matrices. Toeplitz and circulant matrices with their properties (eigenvalue and eigenvector), block matrix computation, and warm-up algorithms.

Unit-2 Introduction to Tensors: Tensors and tensor operations: Mode-n product of a tensor. Kronecker product of two tensors, tensor element product, Khatri-Rao product, the outer product. The Einstein product and t-product tensors. The explicit examples include identity tensor, symmetric tensor, orthogonal tensor, tensor rank, and block tensor.

Unit-3 Tensor Decomposition: Block tensor decomposition, Canonical Polyadic (CP) decomposition, the Tucker decomposition, the multilinear singular value (the higher-order SVD or HOSVD) decomposition, the hierarchical Tucker (HT) decomposition, and the tensor-train (TT) decomposition. Eigenvalue decomposition and singular value decomposition via t-product and the Einstein product. Truncated tensor singular value decomposition. Tensor inversion, and Moore-Penrose inverse. power tensor, solving system of multilinear equations.

Unit-4 Applications of Tensor decompositions: Low-rank tensor approximation, background removal with robust principal tensor component analysis, image deblurring, image compression, compressed sensing with robust Regression, higher-order statistical moments for anomaly detection, solving elliptic partial differential equations.

Unit-5 Tensors for Deep Neural Networks: Deep neural networks, Tensor networks and their decompositions, including, CP decomposition, Tucker decomposition, Hierarchical Tucker decomposition, Tensor train and tensor ring decomposition, Transform-based tensor decomposition. Compressing deep neural networks.

Textbooks / References

- Books
 - Tensors for Data Processing: Theory, Methods, and Applications by Liu, Y. (Ed.), Academic Press (2021)
 - Tensor Computation for Data Analysis by Liu Y, Liu J, Long Z, Zhu C, Springer, 2022
- Recent Articles
 - Block Tensor Unfoldings. *SIAM J. Matrix Anal. Appl.* By S. Ragnarsson and C. F. Van Loan, 33(1):149–169, 2012
 - Orthogonal Tensor Decompositions by setting T. G. Kolda, *SIAM Journal on Matrix Analysis and Applications* 23(1):243–255, 2001
 - Tensor Decompositions and Applications by T. G. Kolda and B. W. Bader, *SIAM Rev.*, 51(3):455–500, 2009
 - An Order-p tensor Factorization with Applications in Imaging by C. D. Martin, R. Shafer, B. Larue, *SIAM J Sci Comput*, 2013;35(1): A474–90
 - Solving Multilinear Systems via Tensor Inversion by M. Brazell, N. Li, C. Navasca, et al, *SIAM J. Matrix Anal Appl.* 2013;34(2):542–570

Prerequisites

Basic linear algebra with basic programming skills (in any programming language)

Grading

- Homework and Quizzes: 50%
- Project: 30%
- End-Term Exam: 20%

E0 251o Data Structures and Graph Analytics (3:1)

Course Instructor

Prof. Viraj Kumar, EECS

Objective

Graph Analytics is important in different domains: Social Networks, Computer Networks, and Computational Biology to name a few. This course deals with the data structures and algorithms underlying Graph Analytics. Important applications will also be considered. Students will be given several programming assignments and a mini project to make them appreciate the contents of the course. Python will be the programming language of implementation for the assignments and the mini project.

Syllabus

Overview of basic data structures such as arrays, linked lists, stacks, and queues. Trees: binary trees and their traversals, binary search trees, balanced trees, and game trees. Priority queues and heaps. Dictionaries and hash tables. Sorting techniques: bubble sort, quicksort, merge sort and decision trees. Algorithm design paradigms: back-tracking, divide and conquer, dynamic programming, and greedy. Graph analytics: depth-first and breadth-first search, connected components and spanning trees, shortest path computation, minimum spanning tree computation, graph matching, network flows, centrality computation, community detection, and connection analysis. Bulk-Synchronous-Parallel model of computation and application to parallel graph algorithms.

Textbooks / References

1. Introduction to Algorithms by T H Cormen, C E Leiserson, and R L Rivest, The MIT Press, Cambridge, Massachusetts, USA, 1990 and 2009 (3rd ed.)
2. Distributed Graph Analytics: Programming, Languages, and their Compilation by Unnikrishnan Cheramangalath, Rupesh Nasre, and Y N Srikant, Springer, 2020
3. Graph Algorithms: Practical Examples in Apache Spark and Neo4j by Amy E, Hodler, and Mark Needham, O'Reilly, 2019 (free book, downloadable)
4. Data Structures and Algorithm Analysis in C++ by Mark A Weiss, Pearson, 4th ed., 2014

Prerequisites

A good knowledge of programming in any programming language. Students will be expected to learn Python on their own.

Grading

- Programming Assignments (2): 10 marks each
- Programming Mini Project (1): 20 marks
- Mid-Term Exams (2): 15 marks each
- End-Term Exam: 30 marks

E0 259o Data Analytics (3:1)

Course Instructors

Prof. Rajesh Sundaresan, ECE

Prof. Ramesh Hariharan, CSA

Prof. Vikram Srinivasan, ECE

Course Description

Data Analytics has assumed increasing importance in recent times. Several industries are now built around the use of data for decision making. Several research areas too, genomics and neuroscience being notable examples, are increasingly focused on large-scale data generation rather than small-scale experimentation to generate initial hypotheses. This brings about a need for data analytics. This course will develop modern statistical tools and modelling techniques through hands-on data analysis in a variety of application domains.

The course will illustrate the principles of hands-on data analytics through several case studies (6 such studies). On each topic, we will introduce a scientific question and discuss why it should be addressed. Next, we will present the available data, how it was collected, etc. We will then discuss models, provide analyses, and finally touch upon how to address the scientific question using the analyses.

We will cover a subset of the following case studies:

1. Astronomy: From Tycho Brahe's observations to the conclusion that Mars moves in an elliptical orbit
2. Sports: The Duckworth-Lewis-Stern method for setting targets in shortened limited overs cricket matches
3. COVID-19: Serological surveys
4. Visual Neuroscience: Neural correlates predict search difficulty
5. Genomics: Understanding the causes of cancer
6. Genomics: The basis for red-green colour blindness
7. Biology: Effects of smoking
8. Networks: Community detection

Syllabus

Data sets from astronomy, genomics, neuroscience, sports, biology, epidemiology, and networks will be analysed to answer specific scientific questions. Statistical tools and modelling techniques will be introduced as needed to analyse the data and eventually address the scientific question. Specific data sets will vary across offerings. Example topics are the following: Tycho Brahe's data on Mars and Kepler's analysis of its orbit (astronomy), the Duckworth-Lewis-Stern method for setting targets in shortened limited overs cricket matches (cricket), retinoblastoma and causes of cancer (genomics), the basis for red-green colour blindness (genomics), serological surveys for COVID-19 (epidemiology), effects of smoking (biology), and community detection (networks)

Textbooks / References

1. Computer Age Statistical Inference by B. Efron and T. Hastie, Cambridge University Press, 2016

Prerequisites

1. Random Processes (E2 202o) OR Probability and Statistics (E0 232) OR equivalent
2. Linear Algebra (E1 219o) OR Matrix Theory

Grading

- Assignments: 40%
- Mid-Term Exam: 30%
- End-Term Exam: 30%

E0 295o Introduction to Cryptography (3:1)

Course Instructors

Prof. Bhavana Kanukurthi, CSA

Prof. Chaya Ganesh, CSA

Course Description

An introductory course in cryptography which will equip the students with the underlying mathematical foundations of the subject ranging from encryption to zero-knowledge.

Syllabus

Modular arithmetic and basic NT, Information theory, one time pad, Shannon impossibility, Computational relaxations, negligible functions, proofs by reduction, PRG intro, PRG defn, next-bit unpred and indistinguishability, equivalence of defns, PRF defn, PRG to PRF GGM, CPA, Cyclic groups, subgroups, structure of \mathbb{Z}_N , \mathbb{Z}_p^* , Assumptions: dlog, DDH, CDH, PRF candidate, Integrity, MACs, CCA defn, construction, Schnorr, Elgamal, RSA, signatures (OTS, FDH), Random oracle model, commitments, Identification schemes, sigma protocols, zero-knowledge proofs

Textbooks / References

1. Introduction to Modern Cryptography by Jonathan Katz and Yehuda Lindell [[Link](#)]
2. A Graduate Course in Applied Cryptography by Dan Boneh and Victor Shoup [[Link](#)]
3. Foundations of Cryptography by Oded Goldreich [[Link](#)]

Prerequisites: Discrete Mathematics, Basic Probability

Grading

HW + Mid-Term Exam + End-Term Exam + Class Participation

E1 277o Reinforcement Learning (3:1)

Course Instructor

Prof. Shalabh Bhatnagar, CSA

Course Description

Reinforcement learning refers to a class of techniques that combine aspects of optimal control, simulation/data driven optimization, and approximation methods for problems of dynamic decision making under uncertainty when the model of the underlying system and its processes is unknown. A large portion of the algorithms and techniques used here are model-free in nature and as a result need no knowledge of the system dynamics and protocols used. Reinforcement Learning thus finds applications in several diverse areas such as Adaptive Control, Signal Processing, Manufacturing, Communication and Wireless Networks, Autonomous Systems and Data Mining.

The objective of this course will be to provide both a rigorous foundation in Reinforcement Learning through the various tools, techniques and algorithms used as well as to cover the state-of-the-art algorithms in Deep Reinforcement Learning involving simulation-based neural network methods.

Syllabus

Introduction to Reinforcement Learning, Multi-armed bandits, Markov decision processes, Dynamic Programming - Value and Policy Iteration Methods, Model-Free Learning Approaches, Monte-Carlo Methods, Temporal Difference Learning, Q-learning, SARSA, Double Q-learning, Value Function Approximation Methods - TD Learning with Linear Function Approximation, Neural Network Architectures, Deep Q-Network Algorithm, Policy Gradient Methods, Actor-Critic Algorithms

Textbooks / References

1. Reinforcement Learning by R. Sutton and A. Barto, MIT Press, 2nd Ed., 2018
2. Reinforcement Learning and Optimal Control by D. Bertsekas, Athena Scientific, 2019
3. Selected Recent Papers

Prerequisites: None

Grading

- Homework: 20%
- Mid-Term Exam: 25%
- Course Project: 25%
- Final-Term Exam: 30%

E1 285o: Advanced Deep Representation Learning (3:1)

Course Instructor

Prof. Prathosh A P, ECE

Course Description

This course contains discussions on cutting-edge topics on Advanced Deep Learning. Given their practical applicability, it is imperative that the algorithms discussed during the lectures are duly implemented (as assignments) by the students. For compute, students can utilize freely-available resources such as Google Colab or Kaggle on toy/small-scale datasets to understand the implementation nuances of different algorithms. There will be 2 assignments and one project component along with terms papers that count for the laboratory credit.

Syllabus

Recap on Fundamentals of Deep Learning: Empirical Risk Minimization, Divergence minimizations and Likelihood maximization Techniques, Deep Learning Architectures (Convolutional and Recurrent Architectures).

Deep Generative Models: Introduction to Generative models, Autoregressive and invertible models, Latent variable models, Variational inference, and recognition networks (VAE, WAE), Adversarial Learning, Generative Adversarial networks, and variants (BiGAN, CycleGAN, StyleGAN, WGAN), Normalizing Flows, Score/Diffusion based models

Transfer Learning and Domain Adaptation: Discrepancy-Based Approaches: statistical (MMD) geometrical and architectural criteria, Generative Domain Adaptation: Adversarial and Non-adversarial Methods, Reconstruction based methods, Domain Generalization: Representation, data manipulation and Learning strategy methods

Few-shot and Meta Learning: Introduction to Multi-task and Transfer learning, Meta-learning framework for few-shot learning, Metric learning, comparators and relational networks, Optimization-based meta learning, Generative meta learning

Semi and Self-supervised Learning: Consistency Regularization, Proxy-label Methods, Active Learning, Weakly supervised learning methods, Self-supervised and Contrastive Representation Learning, Contrastive losses, Memory-bank techniques, BYOL, SWAV, SimCLR, MoCo, Hard negative mining.

Textbooks / References

1. Understanding Machine Learning: From Theory to Algorithms by Shai Ben-David and Shai Shalev-Shwartz, Cambridge University Press
2. Probabilistic Machine Learning: Advanced Topics by Kevin P. Murphy, MIT Press, 2023
3. Deep Learning by Aaron Courville, Ian Goodfellow, and Yoshua Bengio, MIT Press, 2016
4. Mathematics for Machine Learning by Deisenroth, Marc Peter, A. Aldo Faisal, and Cheng Soon Ong, Cambridge University Press, 2020
5. Machine Learning from Weak Supervision: An Empirical Risk Minimization Approach by Masashi Sugiyama, Han Bao, Takashi Ishida, Nan Lu, Tomoya Sakai, and Gang Niu, MIT Press
6. Deep Generative Modeling by Jakub M. Tomczak, Springer 2022
7. Semi-Supervised Learning by Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien, MIT Press
8. Seminal and Survey Papers from Machine Learning Conferences such as ICML, Neurips, ICLR, CVPR, AISTATS etc.

Prerequisites

1. *Mandatory:* A course on probability theory
2. *Mandatory:* A course on classical machine learning fundamentals
3. *Mandatory:* Moderate programming skills in Python

Grading

- 1 Minor and 1 Major: 50%
- Project and Assignments: (20+20) %
- Term Paper: 10%

E1 2860: Deep Generative Models (3:1)

Course Instructor

Prof. Prathosh A P, ECE

Course Description

This graduate-level course provides an in-depth exploration of deep generative models, including their probabilistic foundations and learning algorithms. Students will learn about various types of deep generative models such as variational autoencoders, generative adversarial networks, autoregressive models, normalizing flow models, energy-based models, and score-based models. The course will cover both theoretical foundations and practical implementations of these models using popular frameworks like PyTorch. Students will gain hands-on experience through lectures and assignments, allowing them to explore deep generative models across various AI tasks.

Course Objectives

- Develop a deep understanding of the importance of generative models in artificial intelligence and machine learning
- Design and implement generative models using popular frameworks
- Implement a range of generative models, including autoregressive models, normalizing flow models, energy-based models, and score-based models
- Build problem-solving skills by tackling challenges and complexities in the practical implementation of generative models

Core Competencies

- Deep Generative Models
- Autoregressive Models
- Variational Autoencoders
- Generative Adversarial Networks (GANs)
- Energy-Based Models
- Score-Based Models
- Diffusion Models
- Large Language Models

Syllabus

Recap on Fundamentals of Deep Learning:

Empirical Risk Minimization, Divergence Minimizations and Likelihood Maximization Techniques, Deep Learning Architectures (Convolutional, Recurrent and Transformer Architectures)

Deep Generative Models:

Introduction to Generative Models, Latent Variable Models, Variational Inference and Recognition Networks (Variational Auto Encoders and Wasserstein's Autoencoders)

Adversarial Learning, Generative Adversarial Networks, and variants (BiGAN, CycleGAN, StyleGAN, WGAN)

Diffusion and Score Based Models (DDPM, Consistency Models, DALLE, Stable Diffusion)

Auto-regressive and masked Language Models:

Large Language Models (Encoder/Decoder Models, BERT, GPT and other family)

Multi-Modal Models:

CLIP, Vision-Language Generative Models

Adaptation and Transfer:

Scaling LLMs (Mixture of Experts), Fine Tuning and Adaptation Techniques (LoRA, RAG)

Domain Adversarial Networks, Image-to-Image Translation, and other applications

Textbooks / References

1. Understanding Machine Learning: From Theory to Algorithms by Shai Ben-David and Shai Shalev-Shwartz, Cambridge University Press
2. Probabilistic Machine Learning: Advanced Topics by Kevin P. Murphy, MIT Press, 2023
3. Deep Learning by Aaron Courville, Ian Goodfellow, and Yoshua Bengio, MIT Press, 2016
4. Mathematics for Machine Learning by Deisenroth, Marc Peter, A. Aldo Faisal, and Cheng Soon Ong, Cambridge University Press, 2020
5. Machine Learning from Weak Supervision: An Empirical Risk Minimization Approach by Masashi Sugiyama, Han Bao, Takashi Ishida, Nan Lu, Tomoya Sakai, and Gang Niu, MIT Press
6. Deep Generative Modeling by Jakub M. Tomczak, Springer 2022
7. Deep Learning: Foundations and Concepts, by Christopher Bishop and Hugh Bishop, Springer, 2023
8. Seminal and Survey Papers from Machine Learning Conferences such as ICML, Neurips, ICLR, CVPR, AISTATS etc.

Prerequisites

1. Proficiency in Python: Coding assignments will be in Python.
2. Calculus and Linear Algebra: Students should be comfortable taking (multivariable) derivatives and understand matrix/vector notation and operations.
3. Probability Theory: Students should be familiar with basic probability distributions (Continuous, Gaussian, Bernoulli, etc.) functions.

Grading

- 2 Exams: 20% each (40)
- 3 Programming Assignments: 15% each (45)
- 3 Quizzes: 5% each (15)

E1 294o Edge and Cloud Systems for Machine Learning Algorithms (3:1)

Course Instructor

Prof. Sumit Kumar Mandal, CSA

Course Description

Increasing complexity of machine learning (ML) algorithms have necessitated the emergence of specialized computer systems. Some ML algorithms are executed at edge, some are executed on the cloud. In this course we will delve into different computing kernels of ML algorithms and see how they can be efficiently computed.

Syllabus

- Introduction to Deep Neural Networks (DNN), Convolutional Neural Networks (CNN)
- DNN processing on GPUs
- GPU input pipeline for DNN
- Accelerators for DNNs
- DNN accelerators with emerging technologies
- Accelerators for Graph Convolutional Networks (GCN)
- GCN accelerators with emerging technology
- Accelerators for Large Language Models (LLMs)

The course will involve moderate to heavy programming. The programming part will include building DNNs from scratch, understanding various performance analysis tools in Linux, training ML models with and without GPUs, parallel programming in Python. The programming assignments will be key to understanding ML models from systems perspective.

Textbooks / References

Efficient Processing of Deep Neural Networks by Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, Joel Emer, Morgan & Claypool Publishers

Prerequisites: No prerequisites as such

Grading

- Class Participation: 10%
- Scribe: 10%
- Quiz: 20%
- Homework: 20%
- Project: 40%

E1 3160 Deep Learning for Robotics (2:1)

Course Instructors

Dr. Shishir N. Y. Kolathaya, Dr. Pushpak Jagtap, Dr. Ravi Prakash, RBCCPS

Course Description

This graduate course will explore the new area of interaction between learning and control specifically applied to robotic systems, both from a foundational level together with a view toward application. The course will first build the necessary framework in which to understand basics of robotic systems, including robot kinematics and dynamics, sensing and estimation, control, machine learning, reinforcement learning and imitation learning. With these fundamentals the course will focus on data driven approaches for learning and control applied to robotic systems.

Syllabus

Robot dynamics and kinematics, nonlinear control and stability, Markov decision processes, reinforcement learning, imitation learning, model free methods, exploration vs. exploitation

Note: Active student participation will be expected in course projects, presentations, and delivering lectures in some of the selected topics.

Textbooks / References

1. A Mathematical Introduction to Robotic Manipulation by Richard M. Murray, Zexiang Li, S. Shankar Sastry, CRC Press, 1994
2. Reinforcement Learning: An Introduction by Richard S. Sutton and Andrew G. Barto, MIT Press, 2017
3. Deep Reinforcement Learning by Sergey Levine: <http://rail.eecs.berkeley.edu/deeprlcourse/>
4. Robot Modeling and Control by Mark W. Spong, Seth Hutchinson, M. Vidyasagar, Wiley, 2005

Prerequisites: Students must be well versed with basic mathematical concepts like linear algebra, classical analysis, and probability theory.

Grading

- Homeworks/Quiz: 60%
- Project Reports: 10%
- Project Presentations: 30%

E2 2760 Process Technology and System Engineering (3:1)

Course Instructor

Prof. Hardik Pandya, ESE

Course Description

In many engineering and science areas, miniaturization of the devices presents several challenges, and it requires expertise from several fields, giving rise to interdisciplinary research. A unique opportunity for manufacturing devices within the nm–mm range is possible due to the recent innovations in microfabrication. The six-orders-of-magnitude dimensional span fabricates several devices, including MOSFETs, microfluidics, microheaters, and other microscale sensors, with applications ranging from electronics to computation, clinical research, and robotics. Understanding this research area will also help students think and understand how to design novel devices based on the applications and potentially develop innovative solutions for several existing problems. With a knowledge of microfabrication and systems engineering to integrate the sensors within a data acquisition platform, a student/researcher can work in several fields, from microelectronics, robotics, electronics systems engineering, and instrumentation engineering shows the importance of this course. This course offers exposure to the right mix of concepts in microfabrication for fabricating several microscale electronic devices and their interfacing with data acquisition systems. The course comprehensively explores this exciting field of microengineering technologies, enabling you to understand, design, and build cutting-edge sensor solutions. The course is designed to equip students with the necessary knowledge and skills to fabricate, analyze, and integrate sensors effectively, setting a path to excel in multiple industries.

Syllabus

Introduction and Overview of Microfabrication Process Technology, Overview of Processes used in the Fabrication of Microsensors and Devices, Silicon Wafers - From Sand to the Laboratory, Thin Film Growth and Deposition Techniques, Fundamentals of Material Deposition Techniques: Overview of Physical Vapor Deposition (Sputtering, Thermal and E-beam evaporation) and Chemical Vapor Deposition (LPCVD, PECVD, MOCVD, ALD); Doping - Diffusion and Ion implantation techniques, Lithography – Modes of Lithography, Photoresists: Positive and Negative Photoresists, Masks: Bright Field and Dark Field Masks, Tools in Photolithography, Etching – Wet and Dry Etching, Fabrication of Micro-Engineered Devices: Process Flow of Device Fabrication using Semiconductor Wafers, Data Acquisition Systems Integrated with Signal Conditioning Circuits for Interfacing Sensors and Devices, Case Studies (sensors and transducers)

Lab Component (Videos): Familiarization with Gowning Procedure and Safety Protocols, Introductory Clean Room Visit and Overview of Equipment, Wafer Cleaning Processes, Thermal Evaporation of Metals, E-beam Evaporation of Metals and Insulators, Photolithography, Wet Etching of Metals, Semiconductors and Insulators, Device Fabrication: From Si to Microchips

Textbooks / References

1. Fundamentals of Microfabrication by Marc J Madou
2. Silicon VLSI Technology: Fundamentals, Practice and Modeling by James D. Plummer, Michael D. Deal, Peter D. Griffin
3. Fundamentals of Semiconductor Fabrication by Gary S. May and Simon M. Sze
4. VLSI Fabrication Principles: Silicon and Gallium Arsenide by Sorab K. Ghandhi
5. VLSI Technology by S. M. Sze
6. Fundamentals of Microelectronics by Behzad Razavi
7. Design with Operational Amplifiers and Analog Integrated Circuits by Sergio Franco, 2002, New York: McGraw-Hill
8. Sensors and Signal Conditioning by Ramon Pallás-Areny and John G. Webster, 2012, John Wiley & Sons

Course Outcomes

- Ability to understand the microfabrication process
- Understand the fabrication of PMOS and NMOS transistors
- Understand the fabrication of micro-cantilevers, flexible force sensors, inter-digitated electrodes, and microfluidic platforms
- Design and develop implantable devices for brain-computer interface
- Understand the design criteria and challenges in interfacing micro engineered devices with the real-world using data acquisition systems
- Learn several processes used in microfabrication through experiments in class 1000 and class 10000 cleanroom
- Designing Electronic Modules and acquiring skills in 3D printing

Prerequisites: Undergraduate in Electronics/Physics/Instrumentation/Mechanical

Grading

- Assignments: 20%
- Mid-Term Exam: 30%
- Final-Term Exam: 50%

E9 241o Digital Image Processing (3:1)

Course Instructor

Prof. Chandra Sekhar Seelamantula, EE

Course Description

This is a foundational course in digital image processing and has a significant programming component that augments the concepts taught in the theory classes. This course is essential for all those who wish to pursue advanced research in Computational Imaging, Computer Vision, and Machine Learning using images.

Syllabus

Introduction to image processing; Image acquisition, image representation; Quantization, optimal thresholding, binarization, halftoning; Image histogram, histogram equalization, sharpening; Sampling, aliasing; Fourier transform, magnitude spectrum, phase spectrum, properties; Splines - continuous and discrete image processing; Image filtering - Gaussian smoothing, Bilateral filtering, nonlocal means filtering; Morphological filters; Directional image processing; Image interpolation; Edge detection; Principal component analysis; Pyramidal representation - Gaussian and Laplacian pyramids; Multiscale and multiresolution representation - filter banks and wavelet representations; Image denoising; Image deconvolution; Iterative algorithms for image restoration, wavelet-based algorithms for image restoration; Sparse representations in image processing, dictionary learning; Image super-resolution; Introduction to Radon transform and computerized tomography; Introduction to convolutional neural networks (CNNs) for image processing

Textbooks / References

1. Digital Image Processing by R. Gonzalez and R. Woods, Pearson International, 4th edition, 2018
2. The Essential Guide to Image Processing by A. Bovik, Academic Press, 2009

Prerequisites: None

Grading

- Homework (assigned once in two weeks): 40%
- Mid-Term Exam: 30%
- End-Term Exam: 30%

E9 246o Advanced Topics in Image Processing (3:1)

Course Instructor

Prof. Soma Biswas, EE

Course Description

This course will serve as a smooth transition from the traditional image processing/computer vision techniques to the state-of-the-art deep learning-based techniques. At the end of this course, the students will have a broader perspective of the field and the different real-world applications like object detection, domain adaptation, etc. will help the students understand how to use the techniques taught for real-world applications. The assignments and projects will give them hands-on experience. There will be 3 Lecture hours per week, and coding assignments and projects are a part of this course.

Syllabus

- **Image Features:** Local features, Corner detectors, Scale-Invariant Feature Transform (SIFT) and its variants, Shape detectors like Hough Transform, Texture descriptors like LBP, Evaluation
- **Deep Learning of Image Features:** Convolution layers, fully connected layers, back propagation, non-linear activation, regularization, batch normalization, drop-out, network architectures
- **Image Segmentation:** Graph-based techniques, Active Shape Models, Active Appearance Models, deep learning-based segmentation, evaluation
- **Object Detection:** RCNN and its improved variants like Fast-RCNN, Faster RCNN, YOLO, Evaluation
- **Self-Supervised Learning:** Contrastive Learning, RotNet, etc.
- **Selected Topics:** Vision Transformers for image classification, Multi-modal foundation model like CLIP, few real-world applications, generative models

Textbooks / References

1. Computer Vision: A Modern Approach by David A. Forsyth & Jean Ponce, Prentice Hall
2. Deep Learning by Ian Goodfellow, Yoshua Bengio and Aaron Courville, MIT Press, 2016
3. Recent papers, which will be provided when the topics are taught in class.

Prerequisites: Digital Image Processing (E9 241o)

Grading

- Assignments: 35%
- Mid-Term Exam: 15%
- End-Term Exam: 25%
- Final Project: 25%

DATA SCIENCE AND BUSINESS ANALYTICS
(DSBA)

DSBA Core Courses

DA 201o Applied AI: Building Practical and Scalable ML Systems (3:1)

Course Instructor

Prof. Sashikumar Ganesan, CDS

Course Description

This four-credit course will be offered as a core course for the M.Tech. (Online) DSBA program in place of DA 203o Introduction to Computing for AI & Machine Learning 3:1. The course is designed as a graduate-level (200-series) course.

This course offers a comprehensive introduction to computational thinking with a focus on its practical applications in Artificial Intelligence and Machine Learning (AI & ML). Students will gain a thorough understanding of machine learning techniques and develop skills to design ML systems suitable for production-ready applications.

Syllabus

Python Programming, Parallel Computing, and Machine Learning Tools: Python Environment: Introduction to Python programming, setting up the Python environment, Using Shell commands, Jupyter Notebook, Colab, and VS Code, Python Installation, pip, and virtual environments. Parallel Computing Fundamentals: Understanding computer performance, Caches and Cache optimization techniques, Roofline analysis, Types of parallelism, Shared parallelism, Distributed parallelism, GPUs in Machine Learning, Introduction to GPUs and their architecture, CUDA programming basics. Machine Learning Tools: Pandas, NumPy, Matplotlib, and Seaborn, Understanding shell scripting, Piping commands, awk, sed, and Regular Expressions, Data imputations, Data cleanup with Pandas, Data cleanup with Excel.

Applied Machine Learning: Supervised Learning: Linear Regression, Polynomial Regression, Underfitting and Overfitting, Regularization Methods, Regression Metrics. Classification: Logistic Regression, Loss Functions for Classification, Confusion Matrix, Classification Metrics. Ensemble Classification: Random Forest, Decision Trees, Gini Index, Bagging and Pruning in Random Forest, Boosting, Gradient Boosting, XGBoost. Unsupervised Learning: K-Means Clustering and Elbow Method, Clustering Metrics, Feature Selection and Feature Extraction, Principal Component Analysis (PCA), Singular Value Decomposition (SVD), Linear Discriminant Analysis (LDA), Explained Variance (PCA). Neural Networks: Types of Neural Networks, Single/Multi-layered Perceptron, Activation and Loss Functions, Convolutional Neural Networks, Introduction to Deep Learning.

MLOps: Introduction to MLOps, Overview of MLOps lifecycle. Cloud Infrastructure for MLOps: Cloud services overview, Cloud account setup, Cloud data storage, Compute instance setup. Docker for MLOps: Introduction to Docker, Containerization, Creating Dockerfile for ML model, Building and running Docker container locally, Pushing Docker image to a container registry. ML Model Training and Deployment: Model training and deployment, Model scaling and automation. MLOps Pipelines: Overview of MLOps pipelines, Building an ML pipeline, Adding data validation and error handling to pipeline, Monitoring, and visualizing pipeline. Model Monitoring and Management: Importance of model monitoring and management, Creating model evaluation metrics, Automated model retraining and redeployment. Best Practices for MLOps: Best practices for managing ML models, Security and compliance considerations, Cost optimization strategies, Future trends.

Textbooks

1. Learn Python 3 the hard way: A very simple introduction to the terrifyingly beautiful world of computers and code by Shaw, Zed A, Addison-Wesley Professional, 2017
2. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow by Aurélien Géron, 3rd Edition, O'Reilly Media, Inc. 2022
3. Practical MLOps: Operationalizing Machine Learning Models by Noah Gift & Alfredo Deza, Shroff/O'Reilly; First edition, 2021

Prerequisites: Basic knowledge of programming

Grading

- Homework: 30%
- Mid-Term Exam: 30%
- Final Project: 20%
- End-Term Exam: 20%

DA 204o Data Science in Practice (3:1)

Course Instructor

Prof. Pandarasamy Arjunan, RBCCPS

Course Description

This is a core course for the DSBA stream and an elective for non-DSBA streams. This course provides an introduction to using Data Science, including Machine Learning and Artificial Intelligence, in practice. At the end of the course, students will be confident in identifying and solving end-to-end data science problems in practical applications.

Syllabus

Data Science Fundamentals: Identifying and framing a data science problem in different fields; Data - Types, Pre-processing; Different types of Analytics; Introduction to Machine Learning, Artificial Intelligence; Is ML/AI the right tool for your problem?; Stakeholder Discussion Guidelines; End-to-end Problem Solving through a 6-Step Data Science Process.

Exploratory Data Analysis: Math Foundations of Probability and Statistics, Hypothesis Testing, How much data is sufficient data?; Data Distributions, Imputation, Outlier handling

Data Science for Tabular Data: CART Algorithm, Random Forest Models, Gradient Boosted Models (XGBoost, CatBoost, LightGBM), Feature Importance and Selection, Development-Testing Paradigm, Cross Validation

Deep Learning for Computer Vision and Natural Language Processing: From Linear Regression to Neural Networks from fundamentals, Matrix and Tensor algebra for Neural Networks, Basics of Stochastic Gradient Descent and Backpropagation, Hyperparameter Tuning, Different types of Layers, NN as data-processing pipelines, Practical computer vision with Transfer Learning, Natural Language Processing with Bag of Words models and sequence transformers

Programming for Data Science, ML and AI: Python (NumPy, Pandas), Scikit Learn, TensorFlow, Keras

Textbooks / References

1. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow by Aurélien Géron, 3rd Edition, O'Reilly Media, Inc., 2022
2. A Hands-On Introduction to Data Science by Shah, Chirag, Cambridge University Press, 2020
3. Learn python 3 the hard way: A very simple introduction to the terrifyingly beautiful world of computers and code by Shaw, Zed A, Addison-Wesley Professional, 2017
4. Mathematics for Machine Learning by Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong, Cambridge University Press, 2020 (<https://mml-book.github.io>)

Prerequisites: Basic knowledge of mathematics

Grading

- Programming Assignments: 20%
- Course Project: 20%
- Biweekly Quiz: 30%
- End-Term Exam: 30%

DA 231o Data Engineering at Scale (3:1)

Course Instructor

Prof. Yogesh Simmhan, CDS

Course Description

This four-credit course will be offered every year in the August-December term as a core course of the Data Science and Business Analytics (DSBA) M. Tech (Online) programme. This course is aimed to be an introductory graduate-level (200-series) course. It will motivate the need for Big Data processing, use of distributed systems to design scalable storage and processing systems, and programming and algorithm design on such platforms to implement large-scale data science applications. The course will have several assignments and mini projects. The course lectures will be delivered over video conference and recorded. The students will have access to computing resources on campus or on the cloud that they can access to complete different assignments.

Syllabus

Module 1: Introduction to Big Data storage systems. Motivation for data engineering at scale. Architecture of Google File System/HDFS

Module 2: Introduction to Big Data processing systems. Overview of distributed systems, Cloud computing, strong and weak scaling. Architecture and internals of Apache Spark. Programming using Spark

Module 3: Introduction to relational and NoSQL databases. ACID, BASE and CAP Theorem. Architecture of Dynamo/Cassandra. Overview of Data lakes

Module 4: Introduction to streaming and linked data processing. Architecture and programming of distributed streaming systems like Kafka, Storm and/or Spark Streaming. Architecture and programming of distributed graph processing systems like Pregel/Giraph

Module 5: Machine learning at large scales. Designing ML pipelines. Distributed and federated learning

Module 6: Topics on Big Data and IoT, Cloud Computing, Ethics, etc.

Textbooks / References

1. Hadoop: The Definitive Guide by Tom White
2. Learning Spark by Holden Karau, Patrick Wendell, Matei Zaharia, Andy Konwinski
3. Select reading from literature

Prerequisites: Basics of programming, data structures, algorithms, computer systems

Grading

- Homework and Quizzes: 60%
- Project: 20%
- End-Term Exam: 20%

The course will follow a continuous evaluation philosophy. Students will have short quizzes after each module that will be held during the class and be proctored over video conference. There will be programming assignments and/or mini projects of different modules. A final exam will be held to evaluate their opera learning in the course. This may be a common proctored exam across all courses. Optionally, this may be replaced by a take-home exam.

DSBA Elective Courses

DA 202o Introduction to Data Science (3:1)

Course Instructor

Prof. Deepak Subramani, CDS

Course Description

This course will provide an introduction to Data Science. This four-credit course will be offered every year in the August-December term as a core course to M. Tech (Online) programme. It is aimed to be an introductory graduate-level (200-series) course.

Syllabus

Data Science Fundamentals: Identifying and framing a data science problem in different fields; Data - Types, Pre-processing; Different types of Analytics; Introduction to Machine Learning, Artificial Intelligence

Basic Programming: Data structures, if-else, loops; Visualization; Handling structured data

Probability: Probability axioms, Conditional Probability, Bayes' Theorem, Independence, Counting Problems, Discrete and Continuous Random Variables, Expectation, Iterated Expectation, Total Law of Probability, Covariance, Correlation, Entropy, Mutual Information

Computational Methods: Calculus for Data Science: Functions, Derivative, Partial derivative, Gradient of vector-valued functions and matrices and automatic differentiation, Second derivative Hessian matrix

Linear Algebra: Vectors, Basis, Linear Dependence and Independence, Tensors, Scalars, Inner Products, Outer Product, Norms, Basis, Orthogonal and Orthonormal Vectors, Orthogonalization and Normalization

Matrix Linear Transformation: Frobenius Norm, Matrix Multiplication, Solutions of system of algebraic equations;

Matrix Decomposition: QR Factorization, Singular Value Decomposition; Cholesky Decomposition, Eigenvalue Decomposition

Textbooks / References

1. A Hands-On Introduction to Data Science by Shah, Chirag, Cambridge University Press, 2020
2. Introduction to Probability. Vol. 1. By Dimitri P. Bertsekas, and John N. Tsitsiklis, Belmont, MA, Athena Scientific, 2002
3. Learn python 3 the hard way: A very simple introduction to the terrifyingly beautiful world of computers and code by Zed A. Shaw, Addison-Wesley Professional, 2017
4. Mathematics for Machine Learning by Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong, Cambridge University Press, 2020 (<https://mml-book.github.io>)
5. Linear Algebra and Learning from Data by Gibert Strang, Wellesley-Cambridge Press, 2019
6. Linear Algebra for Everyone by Gibert Strang, Wellesley-Cambridge Press, 2020

Prerequisites: Basic knowledge of mathematics

Grading

- Homework Assignments (four): 40%
- Mid-Term Exam: 30%
- End-Term Exam: 30%

DA 203o Introduction to Computing for AI & Machine Learning (3:1)

Course Instructor

Prof. Sashikumaar Ganesan, CDS

Course Description

This four-credit course will be offered every year in the August-December term as an elective course to M. Tech (Online) programme. This course is aimed to be an introductory graduate-level (200-series) course.

It is intended at building the foundation of computational thinking with applications to Artificial Intelligence and Machine Learning (AI & ML). Besides, how to build a neural network and how to train, evaluate and optimize it with TensorFlow will also be covered in this course.

Syllabus

Programming Foundation: Digital storage of data in computers, memory and data representation, Overflow and Underflow, Round-off errors, the performance of a computer, Caches, Debugging and Profiling, Basic optimization techniques for serial code

Introduction to Python: Object and Data Structure Basics, Python Statements, Methods and Functions, Object-oriented programming (OOP): Inheritance, Encapsulation, Abstraction, Polymorphism. OOP concepts in Python

Python Tools for Data Science: Pandas, NumPy, Matplotlib, Scikit-Learn, Just-in-Time (JIT) compilers, Numba
Computational Thinking: Arrays, Matrix-Vector, Matrix multiplication, Solving dense and sparse systems. Basic machine learning algorithms. Linear Regression, Linear Classification, Multilayer Perceptron, Backpropagation, Automatic Differentiation, Convolutional Networks

Deep Learning with TensorFlow: Tensors, Install TensorFlow, TensorFlow basics, Simple statistics, and plotting, Loading, and exploring data, Learning with TensorFlow and Keras, Mini-project

Textbooks / References

1. Computer Architecture: A Quantitative Approach by John Hennessy David Patterson, 6th edition, Morgan Kaufman, 2017 (<https://www.elsevier.com/books/computer-architecture/hennessy/978-0-12-811905-1>)
2. Learn python 3 the hard way: A very simple introduction to the terrifyingly beautiful world of computers and code by Zed A. Shaw, Addison-Wesley Professional, 2017
3. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow by Aurélien Géron, 2nd Edition, O'Reilly Media, Inc. 2019

Prerequisites: Basic knowledge of mathematics, data structures, and algorithms

Grading

- Homework: 20%
- Mid-Term Exam: 30%
- Final Project: 20%
- End-Term Exam: 30%

DA 212o MLOps at Scale (3:1)

Course Instructor

Prof. Sashikumaar Ganesan, CDS

Course Description

This course is aimed to be an advanced graduate-level (200-series) course. This course is aimed at building the foundation of scalable (parallel) computing for Artificial Intelligence and Machine learning (AI & ML).

Syllabus

Parallel Computer Architecture: Pipelining and super-scalar processor, SIMD vectorization, Caches, Multicore architectures, GPUs, Data access optimization

Programming Models: Shared Memory Programming basics, Shared memory programming with OpenMP, Message-passing, MPI, CUDA, MapReduce

Machine Learning at Scale: Automatic parallelization with Numba, Dask, PySpark, Keras, Distributed training with TensorFlow, Deploying TensorFlow models in AWS

MLOps: Introduction to MLOps, Foundations, MLOps for containers, Continuous Integration, Continuous Deployment for ML models, Monitoring and Feedback

Prerequisites

DA 203o: Introduction to Computing for AI & Machine Learning (or) DA 224o Practical Machine Learning (or) consent of Instructor

Textbooks / References

1. Introduction to High-Performance Computing for Scientists and Engineers by Georg Hager, Gerhard Wellein, CRC Press, 2010
2. Practical MLOps by Noah Gift and Alfredo Deza, 1st Edition, O'Reilly Media, Inc., 2021
3. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow by Aurélien Géron, 2nd Edition, O'Reilly Media, Inc. 2019

Grading

- Assignments: 20%
- Mid-Term Exam: 20%
- Project: 30%
- End-Term Exam: 30%

DA 218o Probabilistic Machine Learning: Theory and Applications (3:1)

Course Instructor

Prof. Punit Rathore, RBCCPS

Course Objective

This course will involve mathematics, practical hands-on workshops in python, and programming tasks that require you to exercise critical thinking. It aims to build a strong foundation in the area of probabilistic/Bayesian techniques for machine learning. I hope that you will find this course and programming tasks intellectually stimulating and relevant to your current program and your careers (research/industry).

Syllabus

Probability Distributions, Concepts on Prior, Likelihoods, and Posteriors Distributions, Maximum Likelihood Estimation, Conjugate Priors, Parameter Estimation
Probabilistic Regression and Classification Models, Regularization using Priors, Ensemble Learning, Model Selection
Probabilistic Unsupervised Learning Models, Latent Variable Models (LVMs) e.g., Gaussian Mixture Models, Expectation-Maximization (EM)
Probabilistic Graphical Models (PGMs), Probabilistic and Statistical Inferencing, Bayesian Estimation, Variable Elimination, Graph Structure Learning, Markov Chain Monte Carlo (MCMC) and Sampling Algorithms, Variational Inference, PGM Examples and Applications, Bayesian Neural Network, Generative Models with Brief Overview on Deep Generative Models, Semi-Supervised Learning, Active Learning, Anomaly Detection Models

Prerequisites

Mandatory: Basic knowledge in Probability (Mandatory), Proficiency in python

Optional: Some basic knowledge in machine learning through coursework or projects

Learning Outcomes

On successful completion of this course, students should be able to

- (i) Gain an understanding of representative selection of ML techniques (not just limited to probabilistic models)
- (ii) Describe and use a range of basic and advanced probabilistic models for data mining and machine learning tasks
- (iii) Code, apply, and solve the real-world problems using python
- (iv) Design, implement, evaluate, and interpret probabilistic models
- (v) Identify limitation of the statistical and probabilistic methods covers in the course
- (vi) Become a discerning ML consumer

Textbooks / References

1. Machine learning: A Probabilistic Perspective by Kevin Murphy, MIT Press, 2012
2. Probabilistic Graphical Models, Principles and Techniques by Daphne Koller, Cambridge University Press, 2009
3. Pattern Recognition and Machine Learning by Christopher Bishop, New York, Springer, 2006
4. Bayesian Reasoning and Machine Learning by David Barber, Cambridge Univ. Press, 2013

Grading

- Two Group Projects/Kaggle: 40%
- Quizzes: 30%
- End-Term Exam: 30%

DA 219o Quantum Computing Methods: Theory and Applications (3:1)

Course Instructor

Prof. Phani Motamarri, CDS

Course Objectives

Quantum computing (QC) has an immense potential to drastically revolutionize industries such as finance, healthcare, AI and automotive over the next several years. This course will give a basic introduction to quantum computing by exposing the participants to the underlying key ideas at the intersection of quantum mechanics, linear algebra, and computing. Additionally, the course will introduce the inner workings of popular quantum algorithms, including how to build, manipulate quantum gates, and circuits using quantum computing toolboxes.

Syllabus

Introduction: Need for quantum computing, quantum vs classical mechanics, quantum vs classical computing, quantum supremacy vs quantum advantage, types of quantum hardware, sneak-peak into a few industrial case studies.

Mathematical Foundations of Quantum Computing: Dirac notation for vectors (Kets, Bra), Linear independence of kets(vectors), basis kets(vectors), linear vector space, orthonormal vectors, inner products, matrices viewed as linear transformations, matrix-vector products, matrix-matrix multiplication, unitary and Hermitian matrices, operators in the quantum world (Hadamard, Pauli X, Y, Z operators), the action of operator on a vector, outer products of vectors, tensor products, the inverse of a matrix, eigenvalue problems, the relevance of these concepts in quantum computing.

Fundamentals of Quantum Computing: Quantum Bit (Qubit), Quantum superposition (qubit state-space), conservation of probabilities, basis transformations, Observables, Time evolution of the quantum system, quantum measurement (eg: spin measurement -via- Stern-Gerlach experiment), single-qubit gates, quantum circuit, n-qubit state space, quantum gates, quantum entanglement, Bell states (EPR states), Bell measurement.

Development Libraries for Quantum Computer Programming: Quantum Computing simulators, Introduction to Qiskit (Quantum Development Kit), Hands-on tutorials for implementation of Quantum Gates, Circuits.

Quantum Algorithms: No-cloning theorem, Quantum algorithm construction, Quantum teleportation, Entanglement swapping, Complexity of algorithms, Deutsch-Jozsa algorithm (1-qubit input, n-qubit input cases), Bernstein-Vazirani algorithm, Grover's algorithm, Introduction to Quantum machine learning, Hands-on sessions for implementation of Quantum algorithms in Qiskit

Textbooks / References

1. Quantum Computing for Everyone by Chris Bernhardt, MIT Press, Cambridge, Massachusetts, 2019
2. Quantum Computing: An Applied Approach by Jack D. Hidary, Springer, 2021 (Second Edition)
3. Quantum Computation and Quantum Information by Michael A Nielsen, Isaac L. Chuang, Cambridge University Press, 2010

Prerequisites: Basic linear Algebra and probability; Some proficiency in Python

Grading

- Homeworks and Projects through Qiskit: 50%
- Mid-Term Exam: 20%
- End-Term Exam: 30%

DA 224o Practical Machine Learning (3:1)

Course Instructor

Prof. Deepak Subramani, CDS

Course Description

This four-credit course will be offered every year in the January-May term as an elective course of the Data Science and Business Analytics (DSBA) Online M. Tech programme. It is aimed to be a graduate-level (200-series) course. The course introduces and trains students in different data-driven modeling approaches and machine learning techniques to succeed in industry and research. Emphasis is laid on both understanding different methods and applying them in practice. At the end of the course, students would be able to use machine learning to model and solve data science problems.

Syllabus

Data-Driven Modelling Concepts: Computational Thinking; Software for Machine Learning: Introduction to Scikit-Learn, Keras and TensorFlow

Supervised Learning: Linear Models for Classification and Regression, Regularization, Optimization Algorithms in Machine Learning, Support Vector Machines (Linear and Kernel), Decision Trees and Ensemble Methods

Dimensionality Reduction: Projection (PCA, kernel PCA), and Manifold Learning (LLE, t-SNE)

Unsupervised Learning: Clustering with K-means, DBSCAN, Gaussian Mixture Models, Anomaly Detection

Basics of Reinforcement Learning: Markov Decision Process, Dynamic Programming, Q-Learning

Bayesian Learning: Bayesian methods in Machine Learning, Basics of Neural Networks and Deep Learning

Textbooks / References

1. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow by Aurélien Géron, O'Reilly Media Inc. (2019)
2. The Elements of Statistical Learning: Data Mining, Inference, and Prediction by Trevor Hastie, Robert Tibshirani, and Jerome Friedman, Springer (2013)
3. Selected chapters, review papers, and online material provided by the instructor

Prerequisites: DA 202o, Introduction to Data Science

Grading

- Homework (Mini Projects): 30%
- Mid-Term Exam: 20%
- Final Project: 20%
- End-Term Exam: 30%

DA 225o Deep Learning (3:1)

Course Instructor

Prof. Deepak Subramani, CDS

Course Description

This four-credit course will be offered every year in the Summer term as an elective course of the Data Science and Business Analytics (DSBA) M. Tech (Online) program. This course is aimed to be a graduate-level (200-series) course. The course introduces and trains students in different deep learning techniques to succeed in industry and research. At the end of the course, students would be able to design and use deep learning algorithms.

Syllabus

Deep Neural Networks, Backpropagation, Regularization and Optimizers; Deep Convolutional Neural Networks; Deep Sequence Modeling with Recurrent Architectures; Generative Modeling with Autoencoders and Generative Adversarial Networks (GANs); Deep Reinforcement Learning; Software for Deep Learning; Selected Applications and Case Studies from Computer Vision, Climate Analytics, Financial Analytics, Autonomous Vehicles

Topics

Module 1: Deep Neural Networks: Building and Training Multilayer Perceptrons, Backpropagation, Gradient Problems, Regularization, Transfer Learning with pre-trained models

Module 2: Deep Convolutional Neural Networks: Convolutional Layers, Pooling Layers, CNN Architectures, Object Detection, Semantic Segmentation. Applications in Computer Vision and Remote Sensing

Module 3: Deep Sequence Modeling with Recurrent Architectures: Recurrent Neurons and Layers, Training Recurrent Neural Networks, Forecasting a Time Series, Handling Long Sequences. Applications in NLP and Financial Analytics

Module 4: Generative Modeling: Stacked Autoencoders, Convolutional Autoencoders, Recurrent Autoencoders, Variational Autoencoders, Generative Adversarial Networks

Module 5: Deep Reinforcement Learning: Policy Gradient Networks, Deep Q Networks. Application in Financial Analytics and Autonomous Underwater Vehicles

Textbooks / References

1. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow by Aurélien Géron, O'Reilly Media Inc. (2019)
2. Deep learning Goodfellow, I., Bengio, Y., & Courville, A., MIT press (2016)
3. Deep learning with Python by Chollet, F., Simon and Schuster (2017)

Prerequisites: DA 224o, Practical Machine Learning or equivalent

Grading

- Homework/Assignments: 40%
- Module Quizzes: 30%
- End-Term Exam: 30%

DA 226o Financial Analytics (3:1)

Course Instructor

Prof. Shashi Jain, Management Studies

Course Description

This four-credit course will be offered as an elective every Aug-Dec term. The course will introduce a number of financial analytics techniques. The course covers the fundamentals of financial time-series, performance models, and forecasting models. The participants will also learn how to evaluate the risk-reward trade-off with an introduction to the Modern Portfolio Theory.

Syllabus

Module 1: Introduction to financial analytics, business forecasting, and time-series data

Module 2: Basic forecasting models, moving averages, exponential smoothing, and Holt-Winter's forecasting model

Module 3: How to identify if a time series is stationary or not and know how to make non-stationary data become stationary, introduction to ARIMA and its implementation

Module 4: Modern portfolio theory and advanced topics on algorithmic trading

Textbooks / References

1. Analysis of Financial Time Series by Ruey Tsay (Wiley Series in Probability and Statistics)
2. Time Series Analysis and its Applications by Shumway and Stoffer
3. Advances in Financial Machine Learning by Marcos Lopez de Prado

Prerequisites: Basics of programming, probability, and statistics

Grading

- Assignments: 30%
- Module Quizzes 30%
- End-Term Exam: 40%

DA 227o Data Mining (3:1)

Course Instructor

Prof. Parthasarathy Ramachandran, Management Studies

Course Description

This is a four-credit elective course that will be offered in the August-December term. The course is intended as an introduction to various data mining techniques. The course will cover Linear models, Linear model selection and Regularization, market basket analysis, classification, and clustering. The course will also give the participants to implement some of the algorithms discussed in the course using the MapReduce framework.

Syllabus

- Introduction to statistical learning, Bias-Variance trade off
- Linear regression, model estimation and assessing the accuracy of the model, Quantitative vs qualitative predictors
- Linear model selection, Shrinkage methods – Ridge regression and Lasso
- Market basket analysis, apriori algorithm, FP-tree construction and projection, association rule interestingness measures
- Classification, Logistic regression, Discriminant analysis, Decision trees - ID3 and C4.5, Bagging, Boosting and Random forests, Naïve Bayes, SVM
- Clustering, K-Means, Mixture models and Expectation Maximization
- Recommendation systems – Content based systems and collaborative filtering
- Mining social networks – clustering social network graphs, communities
- Factor analysis and Principal component analysis
- Error estimation, Resampling methods, k-fold cross validation and bootstrap

Textbooks / References

1. An Introduction to Statistical Learning with Applications in R by Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani
2. Introduction to Data Mining by Pang-Ning Tan, Michael Steinbach, and Vipin Kumar
3. Mining of Massive Datasets by Jure Leskovec, Anand Rajaraman and Jeff Ullman

Prerequisites: Basic programming experience, probability and statistics, linear algebra

Grading

- Assignments: 30
- Final Project: 20
- Mid-Term Exam: 25
- End-Term Exam: 25

DA 245o Linear Optimization and Network Science (3:1)

Course Instructor

Prof. Tarun Rambha, Civil Engineering

Course Description

This course is aimed at:

- (1) formulating practical problems as optimization models and using efficient algorithms for solving them, and
- (2) understanding how dynamic cyber and physical networks evolve.

We will particularly focus on linear and network optimization models which may have continuous or integer variables. These models have many applications in areas such as mobility, scheduling, energy, manufacturing, e-commerce, and logistics. A few hands-on sessions will also introduce you to solvers and visualization tools relevant to the course such as CPLEX/Gurobi, NetworkX, Gephi, and OR tools.

Syllabus

Introduction to Linear Programming (LP); Geometry of LPs; Simplex Method; Duality; Large scale optimization and applications (Column generation, Dantzig Wolfe decomposition, Benders decomposition); Introduction to Networks; Shortest paths (Label setting and label correcting methods, A* algorithm, Contraction hierarchies); Max flows and Min cost problems (Augmenting path method, Cycle cancelling and successive shortest path methods); Integer programs (Branch and bound and cutting plane method); Traveling salesman and Vehicle routing problems; Random networks and centrality (Small worlds, power laws, scale-free properties); Evolution of networks (Preferential attachment); Spreading phenomenon (Epidemics and contact networks).

Textbooks / References

1. Introduction to Linear Optimization (Vol. 6, pp. 479-530) by D. Bertsimas & J.N. Tsitsiklis (1997), Belmont, MA: Athena Scientific
2. Linear Programming and Network Flows by M.S. Bazaraa, J.J. Jarvis & H.D. Sherali (2011), John Wiley & Sons
3. Network Flows: Theory, Algorithms, and Applications by R.K. Ahuja, T.L. Magnanti & J.B. Orlin (1993), Pearson
4. Network Science by A.L. Barabási (2016), Cambridge University Press

Grading

- 5 Quizzes: 30%
- Project: 30%
- End-Term Exam: 40%

DS 216o Applied Artificial Intelligence in Healthcare (3:1)

Course Instructor

Prof. Vaanathi Sundaresan, CDS

Learning Objectives

- To understand the machine learning (ML) concepts including deep learning and choose appropriate techniques/methods for various tasks aiming towards various clinical/healthcare applications.
- To be able to build a data analytics pipeline suited for various real-world applications for various modalities – e.g., audio/speech/sensory signal processing, text data, image analysis applications such as object detection, tracking or counting, etc.
- Evaluate the performance of the method with respect to a gold standard target and analyze the competency of the method.
- Statistical analysis of results (for better analysis of data) for various tasks for given the population/dataset size in the real-world scenarios. This would provide a comprehensive skillset required for tool development, testing and deployment in healthcare industry.

Syllabus

ML Concepts: Introduction to data and types of learning, Expectation maximization methods, clustering; Representation learning; ML classifiers - kernel-based methods, ensemble methods: decision trees, Bayesian networks - hidden Markov models, Conditional random fields, dimensionality reduction, Deep learning – introduction to feedforward networks, feature saliency and visualization, convolutional neural networks, encoder-decoder models, graph-based models, generative models.

ML Applications in Healthcare: Data availability and regulatory framework for AI in healthcare, Template matching, correlation – audio/speech signals; Regression and classification on publicly available speech/sensory/biological signal data, image segmentation, classification & disease prognosis - machine learning classifiers, feature-based and rule-based decision making, image and text correlation, uncertainty estimation, semi-/self-supervised learning.

Evaluation of Analysis Tasks: Evaluation metrics, segmentation evaluation metrics (IoU, Dice, Jaccard indices, Hausdorff distance measures), classification evaluation metrics (confusion matrix, sensitivity, specificity, accuracy), registration metrics (MSE, MAE).

Statistical Evaluation: Testing statistical significance of ML applications: Review of hypothesis testing basic, permutation tests, effect of sample size, statistical power, parametric and non-parametric tests, Shapiro-wilks test, t-tests. Statistical evaluation of ML applications: Descriptive statistics (mean, standard deviation, median, confidence interval, IQR), Unpaired and paired t-tests, one-way and repeated measures ANOVA.

Textbooks / References

1. Pattern Recognition and Machine Learning by C.M. Bishop, Springer, 2006
2. Deep Learning by I. Goodfellow, Y. Bengio and A. Courville, 2016
3. The Elements of Statistical Learning by Jerome H. Friedman, Robert Tibshirani and Trevor Hastie, Springer, 2001
4. Design of Experiments: Statistical Principles of Research Design and Analysis by R.O. Kuehl, 2000
5. Research Papers, Material/Notes provided by instructor.

Prerequisites: Basic knowledge in linear algebra, probability; Proficiency in Python coding

Grading

- Assignments: 15% (best 1 out of 2)
- Mini ML Application Project: 30%
- Mid-Term Exam: 25%,
- End-Term Exam: 30%

DS 261o Artificial Intelligence for Medical Image Analysis (3:1)

Course Instructor

Prof. Phaneendra Yalavarthy, CDS

Learning Outcomes

On successful completion of the course, the student should be able to:

- Identify the basic concepts, terminology, theories, models, and methods for medical image analysis using artificial intelligence
- Characterize the unique challenges associated with various types of medical image data modalities
- Describe and implement the commonly used artificial intelligence methods used for medical image analysis
- Develop and systematically test a number of methods for medical image analysis using artificial intelligence methods
- Choose appropriate evaluation methodologies to evaluate the performance of artificial intelligence methods on biomedical image analysis problems
- Identify limitations of the methods covered in the course

in order to:

- Curate medical image data for use in artificial intelligence-based methods
- Implement, analyze, and evaluate biomedical image analysis systems using artificial intelligence
- Use the knowledge acquired in the course read and profit by literature in the area

Topics

- Overview of biological and medical imaging modalities and research/clinical applications
- Overview of Quick introduction to: (a) medical imaging tools (viewers, formats, etc.) and (b) PyTorch/TensorFlow
- Overview of Basic Mathematics (Linear Algebra, Probability, and Optimization)
- Overview of Challenges in biomedical image data handling and curation
- Overview of Detection / Segmentation / Image classification for biomedical images
- Overview of Machine Learning Methods: SVM, PCA, KNN, FCM, etc. applied to medical image analysis
- Overview of Neural networks and Deep Learning: Principle of learning, CNNs, Loss Functions, etc. for medical image analysis
- Overview of Transfer learning, fine tuning, and generalization in medical image analysis
- Overview of Evaluation methodology in medical image analysis: metrics, calibration, uncertainty, bias, etc.
- Overview of Challenges in healthcare AI deployment: reproducibility, interpretability and regulatory
- Overview of Unsupervised/self-supervised learning in biomedical image analysis
- Overview of Generative models and Inverse problems in medical imaging

Laboratory Component

- Systematically test a number of methods for medical image analysis using artificial intelligence methods
- Mini Project to solve a test problem in medical image analysis (Ex: - Segmentation of Hepatic vessels in kidney using X-ray CT images)

Textbooks / References

1. Probabilistic Machine Learning: An Introduction (Adaptive Computation and Machine Learning series) by Kevin P. Murphy - The MIT Press, March 2022
2. Deep Learning Illustrated: A Visual, Interactive Guide to Artificial Intelligence by Jon Krohn, Grant Beyleveld, and Aglae Bassens – Addison Wesley, 2019

Prerequisites: Basic knowledge of systems and signals, Proficiency in Python

Grading

- Homework: 15% (best 3 out of 4)
- Journal (Review) Paper Presentation: 10%
- Mid-Term Exam: 25%
- Mini Project: 25%
- End-Term Exam: 25%

Guest Lectures by Prasad S. Murthy (Senior Staff AI Scientist, GE HealthCare)

DS 285o Tensor Computations for Data Science (3:1)

Course Instructor

Prof. Ratikanta Behera, CDS

Course Description

This course is an introduction to tensor computations, focusing on theory, algorithms, and applications of tensor decompositions to data sciences. In the era of Big Data, Artificial Intelligence, and Machine Learning, we are faced with the need to process multiway (tensor-shaped) data. These data are mostly in the three or higher order dimensions, whose order of magnitude can reach billions. Huge volumes of multi-dimensional data are a great challenge for processing and analyzing; the matrix representation of data analysis is not enough to represent all the information content of the multiway data in different fields. Further, the importance of being able to decompose a tensor is (at least) two-fold. First, finding the decomposition provides hidden information about the data at hand, and second, having a concise decomposition of the tensor allows us to store it much more efficiently.

The course will provide an understanding of tensor operations and decomposition with various applications, including image deblurring, image compression, neural network, and solving high dimensional partial differential equations.

Learning Outcomes

- Students will learn basic understanding of the theoretical foundations of tensors computation.
- Students will choose efficient tensor decomposition for solving a specific problem.
- Students will learn efficient algorithms for tensor operations; including multiplication, decomposition, and inverse of tensors.
- Students will learn how to implement and use tensor computation in data sciences; including image deblurring, image compression, and solving high dimensional partial differential equations.
- Students will learn the basic tensor computation in the neural networks.

Syllabus

Unit-1 Fundamentals: Basic concepts of matrix properties: norms, rank, trace, inner products, Kronecker product, similarity matrix. Fast Fourier transform, diagonalization of matrices. Toeplitz and circulant matrices with their properties (eigenvalue and eigenvector), block matrix computation, and warm-up algorithms.

Unit-2 Introduction to Tensors: Tensors and tensor operations: Mode-n product of a tensor. Kronecker product of two tensors, tensor element product, Khatri-Rao product, the outer product. The Einstein product and t-product tensors. The explicit examples include identity tensor, symmetric tensor, orthogonal tensor, tensor rank, and block tensor.

Unit-3 Tensor Decomposition: Block tensor decomposition, Canonical Polyadic (CP) decomposition, the Tucker decomposition, the multilinear singular value (the higher-order SVD or HOSVD) decomposition, the hierarchical Tucker (HT) decomposition, and the tensor-train (TT) decomposition. Eigenvalue decomposition and singular value decomposition via t-product and the Einstein product. Truncated tensor singular value decomposition. Tensor inversion, and Moore-Penrose inverse. power tensor, solving system of multilinear equations.

Unit-4 Applications of Tensor Decompositions: Low-rank tensor approximation, background removal with robust principal tensor component analysis, image deblurring, image compression, compressed sensing with robust Regression, higher-order statistical moments for anomaly detection, solving elliptic partial differential equations.

Unit-5 Tensors for Deep Neural Networks: Deep neural networks, Tensor networks and their decompositions, including, CP decomposition, Tucker decomposition, Hierarchical Tucker decomposition, Tensor train and tensor ring decomposition, Transform-based tensor decomposition. Compressing deep neural networks.

Textbooks / References

- Books
 - Tensors for Data Processing: Theory, Methods, and Applications by Liu, Y. (Ed.). Academic Press. (2021)
 - Tensor Computation for Data Analysis by Liu Y, Liu J, Long Z, Zhu C. Springer, 2022
- Recent Articles
 - Block Tensor Unfoldings by S. Ragnarsson and C. F. Van Loan. *SIAM J. Matrix Anal. Appl.*, 33(1):149–169, 2012
 - Orthogonal Tensor Decompositions by T. G. Kolda. *SIAM Journal on Matrix Analysis and Applications*, 23(1):243–255, 2001
 - Tensor Decompositions and Applications by T. G. Kolda and B. W. Bader. *SIAM Rev.*, 51(3):455–500, 2009
 - An Order-p tensor Factorization with Applications in Imaging. *SIAM J Sci Comput.* 2013;35(1): A474–90
 - Solving multilinear systems via tensor inversion by M. Brazell, N. Li, C. Navasca, et al. *SIAM J. Matrix Anal Appl.* 2013;34(2):542–570

Prerequisites

Basic linear algebra with basic programming skills (in any programming language)

Grading

- Homework and Quizzes: 50%
- Project: 30%
- End-Term Exam: 20%

E0 295o Introduction to Cryptography (3:1)

Course Instructors

Prof. Bhavana Kanukurthi, CSA

Prof. Chaya Ganesh, CSA

Course Description

An introductory course in cryptography which will equip the students with the underlying mathematical foundations of the subject ranging from encryption to zero-knowledge.

Syllabus

Modular arithmetic and basic NT, Information theory, one time pad, Shannon impossibility, Computational relaxations, negligible functions, proofs by reduction, PRG intro, PRG defn, next-bit unpred and indistinguishability, equivalence of defns, PRF defn, PRG to PRF GGM, CPA, Cyclic groups, subgroups, structure of \mathbb{Z}_N , \mathbb{Z}_p^* , Assumptions: dlog, DDH, CDH, PRF candidate, Integrity, MACs, CCA defn, construction, Schnorr, Elgamal, RSA, signatures (OTS, FDH), Random oracle model, commitments, Identification schemes, sigma protocols, zero-knowledge proofs

Textbooks / References

1. Introduction to Modern Cryptography by Jonathan Katz and Yehuda Lindell [[Link](#)]
2. A Graduate Course in Applied Cryptography by Dan Boneh and Victor Shoup [[Link](#)]
3. Foundations of Cryptography by Oded Goldreich [[Link](#)]

Prerequisites: Discrete Mathematics, Basic Probability

Grading

HW + Mid-Term Exam + End-Term Exam + Class Participation

E1 220o Linear Algebra (3:1)

Course Instructor

Prof. Sundeep Prabhakar Chepuri, ECE

Course Description

In this course, we will study the basics of linear algebra and matrix theory, with applications to engineering. The focus will be two-fold: on the beautiful mathematical theory of matrices, and their use in solving engineering problems.

Syllabus

Fundamental ideas - vector spaces, matrices, determinant, rank, etc.; Norms, error analysis in linear systems; Eigenvalues and eigenvectors; Canonical, Symmetric and Hermitian forms, matrix factorizations; Least-squares problems, generalized inverses; Miscellaneous topics/applications

Textbooks / References

1. Matrix Analysis by Horn and Johnson, Cambridge University Press
2. Matrix Theory by David Lewis, Allied Publishers
3. Matrix Computations by Golub and Van Loan, 3rd Ed., John Hopkins University Press
4. Linear Algebra and its Applications by Gilbert Strang, 3rd Ed., Harcourth Brace Janovich Pubs. (See also: <http://ocw.mit.edu/courses/mathematics/18-06-linear-algebra-spring-2010/video-lectures/>)

Prerequisites: None

Grading

- Homework Assignments (four): 40%
- Mid-Term Exam: 30%
- End-Term Exam: 30%

E1 252o Linear and Non-Linear Optimization (3:1)

Course Instructor

Prof. Chandramani Singh, ESE

Course Description

In this course, we will study the basics of linear and nonlinear optimization. We will also see several usages of optimization techniques in supervised and unsupervised learning.

Syllabus

Optimization examples, The basics - global vs local optimality; Convex sets, Convex and concave functions; First-order and second-order optimality conditions; Gradient descent methods, Conjugate gradient method, Newton method, Gradient projection method; Constrained optimization with equality and inequality constraints, Duality; Linear programming, simplex method, duality; Barrier and penalty function methods; Sub gradient descent methods; Proximal gradient descent; Augmented Lagrangian methods

Textbooks / References

1. Nonlinear Programming by D. Bertsekas, Athena Scientific, 2016
2. Linear and Nonlinear Programming by D. Luenberger and Y. Ye, Springer, 2008
3. Convex Optimization by S. Boyd and L. Vandenberghe, Cambridge University Press, 2004

Prerequisites: None

Grading

- Homework (assigned approximately once in two weeks): 40%
- Mid-Term Exam: 30%
- End-Term Exam: 30%

Homeworks will be assigned approximately once in two weeks, and a scanned copy of the solutions need to be turned in.

E1 285o: Advanced Deep Representation Learning (3:1)

Course Instructor

Prof. Prathosh A P, ECE

Course Description

This course contains discussions on cutting-edge topics on Advanced Deep learning. Given their practical applicability, it is imperative that the algorithms discussed during the lectures are duly implemented (as assignments) by the students. For compute, students can utilize freely-available resources such as Google Colab or Kaggle on toy/small-scale datasets to understand the implementation nuances of different algorithms. There will be 2 assignments and one project component along with terms papers that count for the laboratory credit.

Syllabus

Recap on Fundamentals of Deep Learning: Empirical Risk Minimization, Divergence minimizations and Likelihood maximization Techniques, Deep Learning Architectures (Convolutional and Recurrent Architectures).

Deep Generative Models: Introduction to Generative models, Autoregressive and invertible models, Latent variable models, Variational inference, and recognition networks (VAE, WAE), Adversarial Learning, Generative Adversarial networks, and variants (BiGAN, CycleGAN, StyleGAN, WGAN), Normalizing Flows, Score/Diffusion based models

Transfer Learning and Domain Adaptation: Discrepancy-Based Approaches: statistical (MMD) geometrical and architectural criteria, Generative Domain Adaptation: Adversarial and Non-adversarial Methods, Reconstruction based methods, Domain Generalization: Representation, data manipulation and Learning strategy methods

Few-shot and Meta Learning: Introduction to Multi-task and Transfer learning, Meta-learning framework for few-shot learning, Metric learning, comparators and relational networks, Optimization-based meta learning, Generative meta learning

Semi and Self-supervised Learning: Consistency Regularization, Proxy-label Methods, Active Learning, Weakly supervised learning methods, Self-supervised and Contrastive Representation Learning, Contrastive losses, Memory-bank techniques, BYOL, SWAV, SimCLR, MoCo, Hard negative mining.

Textbooks / References

1. Understanding Machine Learning: From Theory to Algorithms by Shai Ben-David and Shai Shalev-Shwartz, Cambridge University Press
2. Probabilistic Machine Learning: Advanced Topics by Kevin P. Murphy, MIT Press, 2023
3. Deep Learning by Aaron Courville, Ian Goodfellow, and Yoshua Bengio, MIT Press, 2016
4. Mathematics for Machine Learning by Deisenroth, Marc Peter, A. Aldo Faisal, and Cheng Soon Ong, Cambridge University Press, 2020
5. Machine Learning from Weak Supervision: An Empirical Risk Minimization Approach by Masashi Sugiyama, Han Bao, Takashi Ishida, Nan Lu, Tomoya Sakai, and Gang Niu, MIT Press
6. Deep Generative Modeling by Jakub M. Tomczak, Springer 2022
7. Semi-Supervised Learning by Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien, MIT Press
8. Seminal and Survey Papers from Machine Learning Conferences such as ICML, Neurips, ICLR, CVPR, AISTATS etc.

Prerequisites

1. *Mandatory:* A course on probability theory
2. *Mandatory:* A course on classical machine learning fundamentals
3. *Mandatory:* Moderate programming skills in Python

Grading

- 1 Minor and 1 Major: 50%
- Project and Assignments: (20+20) %
- Term Paper: 10%

E1 2860: Deep Generative Models (3:1)

Course Instructor

Prof. Prathosh A P, ECE

Course Description

This graduate-level course provides an in-depth exploration of deep generative models, including their probabilistic foundations and learning algorithms. Students will learn about various types of deep generative models such as variational autoencoders, generative adversarial networks, autoregressive models, normalizing flow models, energy-based models, and score-based models. The course will cover both theoretical foundations and practical implementations of these models using popular frameworks like PyTorch. Students will gain hands-on experience through lectures and assignments, allowing them to explore deep generative models across various AI tasks.

Course Objectives

- Develop a deep understanding of the importance of generative models in artificial intelligence and machine learning
- Design and implement generative models using popular frameworks
- Implement a range of generative models, including autoregressive models, normalizing flow models, energy-based models, and score-based models
- Build problem-solving skills by tackling challenges and complexities in the practical implementation of generative models

Core Competencies

- Deep Generative Models
- Autoregressive Models
- Variational Autoencoders
- Generative Adversarial Networks (GANs)
- Energy-Based Models
- Score-Based Models
- Diffusion Models
- Large Language Models

Syllabus

Recap on Fundamentals of Deep Learning:

Empirical Risk Minimization, Divergence Minimizations and Likelihood Maximization Techniques, Deep Learning Architectures (Convolutional, Recurrent and Transformer Architectures)

Deep Generative Models:

Introduction to Generative Models, Latent Variable Models, Variational Inference and Recognition Networks (Variational Auto Encoders and Wasserstein's Autoencoders)

Adversarial Learning, Generative Adversarial Networks, and variants (BiGAN, CycleGAN, StyleGAN, WGAN)

Diffusion and Score Based Models (DDPM, Consistency Models, DALLE, Stable Diffusion)

Auto-regressive and masked Language Models:

Large Language Models (Encoder/Decoder Models, BERT, GPT and other family)

Multi-Modal Models:

CLIP, Vision-Language Generative Models

Adaptation and Transfer:

Scaling LLMs (Mixture of Experts), Fine Tuning and Adaptation Techniques (LoRA, RAG)

Domain Adversarial Networks, Image-to-Image Translation, and other applications

Textbooks / References

1. Understanding Machine Learning: From Theory to Algorithms by Shai Ben-David and Shai Shalev-Shwartz, Cambridge University Press
2. Probabilistic Machine Learning: Advanced Topics by Kevin P. Murphy, MIT Press, 2023
3. Deep Learning by Aaron Courville, Ian Goodfellow, and Yoshua Bengio, MIT Press, 2016
4. Mathematics for Machine Learning by Deisenroth, Marc Peter, A. Aldo Faisal, and Cheng Soon Ong, Cambridge University Press, 2020
5. Machine Learning from Weak Supervision: An Empirical Risk Minimization Approach by Masashi Sugiyama, Han Bao, Takashi Ishida, Nan Lu, Tomoya Sakai, and Gang Niu, MIT Press
6. Deep Generative Modeling by Jakub M. Tomczak, Springer 2022
7. Deep Learning: Foundations and Concepts, by Christopher Bishop and Hugh Bishop, Springer, 2023
8. Seminal and Survey Papers from Machine Learning Conferences such as ICML, Neurips, ICLR, CVPR, AISTATS etc.

Prerequisites

1. Proficiency in Python: Coding assignments will be in Python.
2. Calculus and Linear Algebra: Students should be comfortable taking (multivariable) derivatives and understand matrix/vector notation and operations.
3. Probability Theory: Students should be familiar with basic probability distributions (Continuous, Gaussian, Bernoulli, etc.) functions.

Grading

- 2 Exams: 20% each (40)
- 3 Programming Assignments: 15% each (45)
- 3 Quizzes: 5% each (15)

E1 294o Edge and Cloud Systems for Machine Learning Algorithms (3:1)

Course Instructor

Prof. Sumit Kumar Mandal, CSA

Course Description

Increasing complexity of machine learning (ML) algorithms have necessitated the emergence of specialized computer systems. Some ML algorithms are executed at edge, some are executed on the cloud. In this course we will delve into different computing kernels of ML algorithms and see how they can be efficiently computed.

Syllabus

- Introduction to Deep Neural Networks (DNN), Convolutional Neural Networks (CNN)
- DNN processing on GPUs
- GPU input pipeline for DNN
- Accelerators for DNNs
- DNN accelerators with emerging technologies
- Accelerators for Graph Convolutional Networks (GCN)
- GCN accelerators with emerging technology
- Accelerators for Large Language Models (LLMs)

The course will involve moderate to heavy programming. The programming part will include building DNNs from scratch, understanding various performance analysis tools in Linux, training ML models with and without GPUs, parallel programming in Python. The programming assignments will be key to understanding ML models from systems perspective.

Textbooks / References

Efficient Processing of Deep Neural Networks by Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, Joel Emer, Morgan & Claypool Publishers

Prerequisites: No prerequisites as such

Grading

- Class Participation: 10%
- Scribe: 10%
- Quiz: 20%
- Homework: 20%
- Project: 40%

E1 3160 Deep Learning for Robotics (2:1)

Course Instructors

Dr. Shishir N. Y. Kolathaya, Dr. Pushpak Jagtap, Dr. Ravi Prakash, RBCCPS

Course Description

This graduate course will explore the new area of interaction between learning and control specifically applied to robotic systems, both from a foundational level together with a view toward application. The course will first build the necessary framework in which to understand basics of robotic systems, including robot kinematics and dynamics, sensing and estimation, control, machine learning, reinforcement learning and imitation learning. With these fundamentals the course will focus on data driven approaches for learning and control applied to robotic systems.

Syllabus

Robot dynamics and kinematics, nonlinear control and stability, Markov decision processes, reinforcement learning, imitation learning, model free methods, exploration vs. exploitation

Note: Active student participation will be expected in course projects, presentations, and delivering lectures in some of the selected topics.

Textbooks / References

1. A Mathematical Introduction to Robotic Manipulation by Richard M. Murray, Zexiang Li, S. Shankar Sastry, CRC Press, 1994
2. Reinforcement Learning: An Introduction by Richard S. Sutton and Andrew G. Barto, MIT Press, 2017
3. Deep Reinforcement Learning by Sergey Levine: <http://rail.eecs.berkeley.edu/deeprlcourse/>
4. Robot Modeling and Control by Mark W. Spong, Seth Hutchinson, M. Vidyasagar, Wiley, 2005

Prerequisites: Students must be well versed with basic mathematical concepts like linear algebra, classical analysis, and probability theory.

Grading

- Homeworks/Quiz: 60%
- Project Reports: 10%
- Project Presentations: 30%

E2 2760 Process Technology and System Engineering (3:1)

Course Instructor

Prof. Hardik Pandya, ESE

Course Description

In many engineering and science areas, miniaturization of the devices presents several challenges, and it requires expertise from several fields, giving rise to interdisciplinary research. A unique opportunity for manufacturing devices within the nm–mm range is possible due to the recent innovations in microfabrication. The six-orders-of-magnitude dimensional span fabricates several devices, including MOSFETs, microfluidics, microheaters, and other microscale sensors, with applications ranging from electronics to computation, clinical research, and robotics. Understanding this research area will also help students think and understand how to design novel devices based on the applications and potentially develop innovative solutions for several existing problems. With a knowledge of microfabrication and systems engineering to integrate the sensors within a data acquisition platform, a student/researcher can work in several fields, from microelectronics, robotics, electronics systems engineering, and instrumentation engineering shows the importance of this course. This course offers exposure to the right mix of concepts in microfabrication for fabricating several microscale electronic devices and their interfacing with data acquisition systems. The course comprehensively explores this exciting field of microengineering technologies, enabling you to understand, design, and build cutting-edge sensor solutions. The course is designed to equip students with the necessary knowledge and skills to fabricate, analyze, and integrate sensors effectively, setting a path to excel in multiple industries.

Syllabus

Introduction and Overview of Microfabrication Process Technology, Overview of Processes used in the Fabrication of Microsensors and Devices, Silicon Wafers - From Sand to the Laboratory, Thin Film Growth and Deposition Techniques, Fundamentals of Material Deposition Techniques: Overview of Physical Vapor Deposition (Sputtering, Thermal and E-beam evaporation) and Chemical Vapor Deposition (LPCVD, PECVD, MOCVD, ALD); Doping - Diffusion and Ion implantation techniques, Lithography – Modes of Lithography, Photoresists: Positive and Negative Photoresists, Masks: Bright Field and Dark Field Masks, Tools in Photolithography, Etching – Wet and Dry Etching, Fabrication of Micro-Engineered Devices: Process Flow of Device Fabrication using Semiconductor Wafers, Data Acquisition Systems Integrated with Signal Conditioning Circuits for Interfacing Sensors and Devices, Case Studies (sensors and transducers)

Lab Component (Videos): Familiarization with Gowning Procedure and Safety Protocols, Introductory Clean Room Visit and Overview of Equipment, Wafer Cleaning Processes, Thermal Evaporation of Metals, E-beam Evaporation of Metals and Insulators, Photolithography, Wet Etching of Metals, Semiconductors and Insulators, Device Fabrication: From Si to Microchips

Textbooks / References

1. Fundamentals of Microfabrication by Marc J Madou
2. Silicon VLSI Technology: Fundamentals, Practice and Modeling by James D. Plummer, Michael D. Deal, Peter D. Griffin
3. Fundamentals of Semiconductor Fabrication by Gary S. May and Simon M. Sze
4. VLSI Fabrication Principles: Silicon and Gallium Arsenide by Sorab K. Ghandhi
5. VLSI Technology by S. M. Sze
6. Fundamentals of Microelectronics by Behzad Razavi
7. Design with Operational Amplifiers and Analog Integrated Circuits by Sergio Franco, 2002, New York: McGraw-Hill
8. Sensors and Signal Conditioning by Ramon Pallás-Areny and John G. Webster, 2012, John Wiley & Sons

Course Outcomes

- Ability to understand the microfabrication process
- Understand the fabrication of PMOS and NMOS transistors
- Understand the fabrication of micro-cantilevers, flexible force sensors, inter-digitated electrodes, and microfluidic platforms
- Design and develop implantable devices for brain-computer interface
- Understand the design criteria and challenges in interfacing micro engineered devices with the real-world using data acquisition systems
- Learn several processes used in microfabrication through experiments in class 1000 and class 10000 cleanroom
- Designing Electronic Modules and acquiring skills in 3D printing

Prerequisites: Undergraduate in Electronics/Physics/Instrumentation/Mechanical

Grading

- Assignments: 20%
- Mid-Term Exam: 30%
- Final-term Exam: 50%

E9 2460 Advanced Topics in Image Processing (3:1)

Course Instructor

Prof. Soma Biswas, EE

Course Description

This course will serve as a smooth transition from the traditional image processing/computer vision techniques to the state-of-the-art deep learning-based techniques. At the end of this course, the students will have a broader perspective of the field and the different real-world applications like object detection, domain adaptation, etc. will help the students understand how to use the techniques taught for real-world applications. The assignments and projects will give them hands-on experience. There will be 3 Lecture hours per week, and coding assignments and projects are a part of this course.

Syllabus

- **Image Features:** Local features, Corner detectors, Scale-Invariant Feature Transform (SIFT) and its variants, Shape detectors like Hough Transform, Texture descriptors like LBP, Evaluation
- **Deep Learning of Image Features:** Convolution layers, fully connected layers, back propagation, non-linear activation, regularization, batch normalization, drop-out, network architectures
- **Image Segmentation:** Graph-based techniques, Active Shape Models, Active Appearance Models, deep learning-based segmentation, evaluation
- **Object Detection:** RCNN and its improved variants like Fast-RCNN, Faster RCNN, YOLO, Evaluation
- **Self-Supervised Learning:** Contrastive Learning, RotNet, etc.
- **Selected Topics:** Vision Transformers for image classification, Multi-modal foundation model like CLIP, few real-world applications, generative models

Textbooks / References

4. Computer Vision: A Modern Approach by David A. Forsyth & Jean Ponce, Prentice Hall
5. Deep Learning by Ian Goodfellow, Yoshua Bengio and Aaron Courville, MIT Press, 2016
6. Recent papers, which will be provided when the topics are taught in class.

Prerequisites: Digital Image Processing (E9 241o)

Grading

- Assignments: 35%
- Mid-Term Exam: 15%
- End-Term Exam: 25%
- Final Project: 25%

ELECTRONICS AND COMMUNICATION
ENGINEERING
(ECE)

ECE Core Courses

E1 220o Linear Algebra (3:1)

Course Instructor

Prof. Sundeep Prabhakar Chepuri, ECE

Course Description

In this course, we will study the basics of linear algebra and matrix theory, with applications to engineering. The focus will be two-fold: on the beautiful mathematical theory of matrices, and their use in solving engineering problems.

Syllabus

Fundamental ideas - vector spaces, matrices, determinant, rank, etc.; Norms, error analysis in linear systems; Eigenvalues and eigenvectors; Canonical, Symmetric and hermitian forms, matrix factorizations; Least-squares problems, generalized inverses; Miscellaneous topics/applications

Textbooks / References

1. Matrix Analysis by Horn and Johnson, Cambridge University Press
2. Matrix Theory by David Lewis, Allied Publishers
3. Matrix Computations by Golub and Van Loan, 3rd Ed., John Hopkins University Press
4. Linear Algebra and its Applications by Gilbert Strang, 3rd Ed., Harcourth Brace Janovich Pubs (See also: <http://ocw.mit.edu/courses/mathematics/18-06-linear-algebra-spring-2010/video-lectures/>)

Prerequisites: None

Grading

- Homework Assignments (four): 40%
- Mid-Term Exam: 30%
- End-Term Exam: 30%

E1 245o Statistical Inference for Engineers and Data Scientists (3:1)

Course Instructor

Prof. Vaibhav Katewa, ECE

Course Description

The main goal of this course is to cover the major domains of statistical inference, namely, estimation, detection, and learning, which include the many mathematical tools that engineers and statisticians use to draw inference from imperfect or incomplete data. The first part of the course develops statistical parameter estimation methods to extract information from data in noise. The second part of this course is about the application of statistical hypothesis testing to the detection of data in noise.

Syllabus

Minimum variance unbiased estimators, maximum likelihood theory, the Cramer-Rao bound, best linear unbiased estimators, Bayesian estimation techniques, the Wiener and Kalman filters, generative models, regression and regularization, binary and multiple hypothesis testing, Neyman-Pearson detector, Bayes detector, and composite hypothesis testing

Course Level

This is a 3:1 credit core course for the M. Tech (Online) ECE program and relevant for AI and DSBA programs as well.

Textbooks / References

1. Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory by S.M. Kay, Prentice Hall 1993, ISBN-13: 978-0133457117
2. Fundamentals of Statistical Signal Processing, Volume II: Detection Theory by S.M. Kay, Prentice 1993, ISBN-13: 978-0135041352

Prerequisites: Matrix theory and Random processes

Grading and Other Requirements

- Three Assignments: 10% each (i.e., 30 % in total)
- Mid-Term Exam: 20%
- Final Project: 30%
- End-Term Exam: 20%

The three assignments and the mid-term exam shall be treated as sessional assessment, while the final project and final-term exam shall be treated as the final assessment.

E2 2020 Random Processes (3:1)

Course Instructor

Prof. Aditya Gopalan, ECE

Course Description

This course is a graduate-level course on probability and stochastic processes. It is assumed that the students are familiar with multivariable calculus (functions of several variables, partial derivatives, integration in n-dimensional real Euclidean spaces) and that they have some idea of elementary probability (e.g., as part of Foundations for Business Analytics of an undergraduate course on mathematics). Some familiarity with vector spaces and matrices would be assumed. The course would be useful for first year Masters or Ph.D. students and would equip them with basic background in probability which is required for more advanced courses such as Machine learning, Adaptive Signal Processing etc. The course is a mathematics course, and the students are encouraged to solve many problems. There would be some tutorial classes to help students with problem solving.

Syllabus

The axioms of probability theory, probability spaces, conditional probability, independence, random variables and distribution functions, continuous and discrete random variables, multiple random variables and joint distributions, conditional distributions, functions of random variables and random vectors, expectation and moments, conditional expectation, some moment inequalities, sequences of random variables and convergence concepts, laws of large numbers, sums of independent random variables and the central limit theorem, stochastic processes, stationarity and ergodicity, discrete time Markov chains, Poisson process, continuous time Markov chains, Brownian motion

Textbooks / References

1. An Introduction to Probability and Statistics by V. K. Rohatgi and A. K. M. E. Saleh
2. Introduction to Probability Theory by P. G. Hoel, S. C. Port and C. J. Stone
3. Introduction to Stochastic Processes by P. G. Hoel, S. C. Port and C. J. Stone
4. Introduction to Probability Models by S. M. Ross

Prerequisites: There are no formal prerequisites. However, students should be familiar with multivariate calculus.

Grading

- Homework Assignments (four): 40%
- Mid-Term Exam: 30%
- End-Term Exam: 30%

ECE Elective Courses

DA 201o Applied AI: Building Practical and Scalable ML Systems (3:1)

Course Instructor

Prof. Sashikumaar Ganesan, CDS

Course Description

This four-credit course will be offered as a core course for the M.Tech. (Online) DSBA program in place of DA 203o Introduction to Computing for AI & Machine Learning 3:1. The course is designed as a graduate-level (200-series) course.

This course offers a comprehensive introduction to computational thinking with a focus on its practical applications in Artificial Intelligence and Machine Learning (AI & ML). Students will gain a thorough understanding of machine learning techniques and develop skills to design ML systems suitable for production-ready applications.

Syllabus

Python Programming, Parallel Computing, and Machine Learning Tools: Python Environment: Introduction to Python programming, Setting up the Python environment, Using Shell commands, Jupyter Notebook, Colab, and VS Code, Python Installation, pip, and virtual environments. Parallel Computing Fundamentals: Understanding computer performance, Caches and Cache optimization techniques, Roofline analysis, Types of parallelism, Shared parallelism, Distributed parallelism, GPUs in Machine Learning, Introduction to GPUs and their architecture, CUDA programming basics. Machine Learning Tools: Pandas, NumPy, Matplotlib, and Seaborn, Understanding shell scripting, Piping commands, awk, sed, and Regular Expressions, Data imputations, Data cleanup with Pandas, Data cleanup with Excel.

Applied Machine Learning: Supervised Learning: Linear Regression, Polynomial Regression, Underfitting and Overfitting, Regularization Methods, Regression Metrics. Classification: Logistic Regression, Loss Functions for Classification, Confusion Matrix, Classification Metrics. Ensemble Classification: Random Forest, Decision Trees, Gini Index, Bagging and Pruning in Random Forest, Boosting, Gradient Boosting, XGBoost. Unsupervised Learning: K-Means Clustering and Elbow Method, Clustering Metrics, Feature Selection and Feature Extraction, Principal Component Analysis (PCA), Singular Value Decomposition (SVD), Linear Discriminant Analysis (LDA), Explained Variance (PCA). Neural Networks: Types of Neural Networks, Single/Multi-layered Perceptron, Activation and Loss Functions, Convolutional Neural Networks, Introduction to Deep Learning.

MLOps: Introduction to MLOps, Overview of MLOps lifecycle. Cloud Infrastructure for MLOps: Cloud services overview, Cloud account setup, Cloud data storage, Compute instance setup. Docker for MLOps: Introduction to Docker, Containerization, Creating Dockerfile for ML model, Building and running Docker container locally, Pushing Docker image to a container registry. ML Model Training and Deployment: Model training and deployment, Model scaling and automation. MLOps Pipelines: Overview of MLOps pipelines, Building an ML pipeline, Adding data validation and error handling to pipeline, Monitoring, and visualizing pipeline. Model Monitoring and Management: Importance of model monitoring and management, Creating model evaluation metrics, Automated model retraining and redeployment. Best Practices for MLOps: Best practices for managing ML models, Security and compliance considerations, Cost optimization strategies, Future trends.

Textbooks / References

1. Learn Python 3 the hard way: A very simple introduction to the terrifyingly beautiful world of computers and code by Shaw, Zed A, Addison-Wesley Professional, 2017
2. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow by Aurélien Géron, 3rd Edition, O'Reilly Media, Inc. 2022
3. Practical MLOps: Operationalizing Machine Learning Models by Noah Gift & Alfredo Deza, Shroff/O'Reilly; First edition, 2021

Prerequisites: Basic knowledge of programming

Grading

- Homework: 30%
- Midterm: 30%
- Final Project: 20%
- End-Term Exam: 20%

DA 204o Data Science in Practice (3:1)

Course Instructor

Prof. Pandarasamy Arjunan, RBCCPS

Course Description

This is a core course for the DSBA stream and an elective for non-DSBA streams. This course provides an introduction to using Data Science, including Machine Learning and Artificial Intelligence, in practice. At the end of the course, students will be confident in identifying and solving end-to-end data science problems in practical applications.

Syllabus

Data Science Fundamentals: Identifying and framing a data science problem in different fields; Data - Types, Pre-processing; Different types of Analytics; Introduction to Machine Learning, Artificial Intelligence; Is ML/AI the right tool for your problem?; Stakeholder Discussion Guidelines; End-to-end Problem Solving through a 6-Step Data Science Process.

Exploratory Data Analysis: Math Foundations of Probability and Statistics, Hypothesis Testing, How much data is sufficient data?, Data Distributions, Imputation, Outlier handling

Data Science for Tabular Data: CART Algorithm, Random Forest Models, Gradient Boosted Models (XGBoost, CatBoost, LightGBM), Feature Importance and Selection, Development-Testing Paradigm, Cross Validation

Deep Learning for Computer Vision and Natural Language Processing: From Linear Regression to Neural Networks from fundamentals, Matrix and Tensor algebra for Neural Networks, Basics of Stochastic Gradient Descent and Backpropagation, Hyperparameter Tuning, Different types of Layers, NN as data-processing pipelines, Practical computer vision with Transfer Learning, Natural Language Processing with Bag of Words models and sequence transformers

Programming for data science, ML and AI: Python (NumPy, Pandas), Scikit Learn, TensorFlow, Keras

Textbooks / References

1. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow by Aurélien Géron, 3rd Edition, O'Reilly Media, Inc., 2022
2. A Hands-On Introduction to Data Science by Shah, Chirag, Cambridge University Press, 2020
3. Learn python 3 the hard way: A very simple introduction to the terrifyingly beautiful world of computers and code by Shaw, Zed A, Addison-Wesley Professional, 2017
4. Mathematics for Machine Learning by Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong, Cambridge University Press, 2020 (<https://mml-book.github.io>)

Prerequisites: Basic knowledge of mathematics

Grading

- Programming Assignments: 20%
- Course Project: 20%
- Biweekly Quiz: 30%
- End-Term Exam: 30%

DA 219o Quantum Computing Methods: Theory and Applications (3:1)

Course Instructor

Prof. Phani Motamarri, CDS

Course Objectives

Quantum computing (QC) has an immense potential to drastically revolutionize industries such as finance, healthcare, AI and automotive over the next several years. This course will give a basic introduction to quantum computing by exposing the participants to the underlying key ideas at the intersection of quantum mechanics, linear algebra, and computing. Additionally, the course will introduce the inner workings of popular quantum algorithms, including how to build, manipulate quantum gates, and circuits using quantum computing toolboxes.

Syllabus

Introduction: Need for quantum computing, quantum vs classical mechanics, quantum vs classical computing, quantum supremacy vs quantum advantage, types of quantum hardware, sneak-peak into a few industrial case studies.

Mathematical Foundations of Quantum Computing: Dirac notation for vectors (Kets, Bra), Linear independence of kets (vectors), basis kets (vectors), linear vector space, orthonormal vectors, inner products, matrices viewed as linear transformations, matrix-vector products, matrix-matrix multiplication, unitary and Hermitian matrices, operators in the quantum world (Hadamard, Pauli X, Y, Z operators), the action of operator on a vector, outer products of vectors, tensor products, the inverse of a matrix, eigenvalue problems, the relevance of these concepts in quantum computing.

Fundamentals of Quantum Computing: Quantum Bit (Qubit), Quantum superposition (qubit state-space), conservation of probabilities, basis transformations, Observables, Time evolution of the quantum system, quantum measurement (eg: spin measurement -via- Stern-Gerlach experiment), single-qubit gates, quantum circuit, n-qubit state space, quantum gates, quantum entanglement, Bell states (EPR states), Bell measurement.

Development Libraries for Quantum Computer Programming: Quantum Computing simulators, Introduction to Qiskit (Quantum Development Kit), Hands-on tutorials for implementation of Quantum Gates, Circuits.

Quantum Algorithms: No-cloning theorem, Quantum algorithm construction, Quantum teleportation, Entanglement swapping, Complexity of algorithms, Deutsch-Jozsa algorithm (1-qubit input, n-qubit input cases), Bernstein-Vazirani algorithm, Grover's algorithm, Introduction to Quantum machine learning, Hands-on sessions for implementation of Quantum algorithms in Qiskit

Textbooks / References

1. Quantum Computing for Everyone by Chris Bernhardt, MIT Press, Cambridge, Massachusetts, 2019
2. Quantum Computing: An Applied Approach by Jack D. Hidary, Springer, 2021 (Second Edition)
3. Quantum Computation and Quantum Information by Michael A Nielsen, Isaac L. Chuang, Cambridge University Press, 2010

Prerequisites: Basic Linear algebra and Probability; Some proficiency in Python

Grading

- Homeworks and Projects through Qiskit: 50%
- Mid-Term Exam: 20%
- End-Term Exam: 30%

DA 245o Linear Optimization and Network Science (3:1)

Course Instructor

Prof. Tarun Rambha, Civil Engineering

Course Description

This course is aimed at:

- (1) formulating practical problems as optimization models and using efficient algorithms for solving them, and
- (2) understanding how dynamic cyber and physical networks evolve.

We will particularly focus on linear and network optimization models which may have continuous or integer variables. These models have many applications in areas such as mobility, scheduling, energy, manufacturing, e-commerce, and logistics. A few hands-on sessions will also introduce you to solvers and visualization tools relevant to the course such as CPLEX/Gurobi, NetworkX, Gephi, and OR tools.

Syllabus

Introduction to Linear Programming (LP); Geometry of LPs; Simplex Method; Duality; Large scale optimization and applications (Column generation, Dantzig Wolfe decomposition, Benders decomposition); Introduction to Networks; Shortest paths (Label setting and label correcting methods, A* algorithm, Contraction hierarchies); Max flows and Min cost problems (Augmenting path method, Cycle cancelling and successive shortest path methods); Integer programs (Branch and bound and cutting plane method); Traveling salesman and Vehicle routing problems; Random networks and centrality (Small worlds, power laws, scale-free properties); Evolution of networks (Preferential attachment); Spreading phenomenon (Epidemics and contact networks).

Textbooks / References

1. Introduction to Linear Optimization (Vol. 6, pp. 479-530) by D. Bertsimas & J.N. Tsitsiklis (1997), Belmont, MA: Athena Scientific
2. Linear Programming and Network Flows by M.S. Bazaraa, J.J. Jarvis & H.D. Sherali (2011), John Wiley & Sons
3. Network Flows: Theory, Algorithms, and Applications by R.K. Ahuja, T.L. Magnanti & J.B. Orlin (1993), Pearson
4. Network Science by A.L. Barabási (2016), Cambridge University Press

Grading

- 5 Quizzes: 30%
- Project: 30%
- End-Term Exam: 40%

DS 216o Applied Artificial Intelligence in Healthcare (3:1)

Course Instructor

Prof. Vaanathi Sundaresan, CDS

Learning Objectives

- To understand the machine learning (ML) concepts including deep learning and choose appropriate techniques/methods for various tasks aiming towards various clinical/healthcare applications.
- To be able to build a data analytics pipeline suited for various real-world applications for various modalities – e.g., audio/speech/sensory signal processing, text data, image analysis applications such as object detection, tracking or counting, etc.
- Evaluate the performance of the method with respect to a gold standard target and analyze the competency of the method.
- Statistical analysis of results (for better analysis of data) for various tasks for given the population/dataset size in the real-world scenarios. This would provide a comprehensive skillset required for tool development, testing and deployment in healthcare industry.

Syllabus

ML Concepts: Introduction to data and types of learning, Expectation maximization methods, clustering; Representation learning; ML classifiers - kernel-based methods, ensemble methods: decision trees, Bayesian networks - hidden Markov models, Conditional random fields, dimensionality reduction, Deep learning – introduction to feedforward networks, feature saliency and visualization, convolutional neural networks, encoder-decoder models, graph-based models, generative models.

ML Applications in Healthcare: Data availability and regulatory framework for AI in healthcare, Template matching, correlation – audio/speech signals; Regression and classification on publicly available speech/sensory/biological signal data, image segmentation, classification & disease prognosis - machine learning classifiers, feature-based and rule-based decision making, image and text correlation, uncertainty estimation, semi-/self-supervised learning.

Evaluation of Analysis Tasks: Evaluation metrics, segmentation evaluation metrics (IoU, Dice, Jaccard indices, Hausdorff distance measures), classification evaluation metrics (confusion matrix, sensitivity, specificity, accuracy), registration metrics (MSE, MAE).

Statistical Evaluation: Testing statistical significance of ML applications: Review of hypothesis testing basic, permutation tests, effect of sample size, statistical power, parametric and non-parametric tests, Shapiro-wilks test, t-tests. Statistical evaluation of ML applications: Descriptive statistics (mean, standard deviation, median, confidence interval, IQR), Unpaired and paired t-tests, one-way and repeated measures ANOVA.

Textbooks / References

1. Pattern Recognition and Machine Learning by C.M. Bishop, Springer, 2006
2. Deep Learning by I. Goodfellow, Y. Bengio and A. Courville, 2016
3. The Elements of Statistical Learning by Jerome H. Friedman, Robert Tibshirani and Trevor Hastie, Springer, 2001
4. Design of Experiments: Statistical Principles of Research Design and Analysis by RO Kuehl, 2000
5. Research Papers, Material/Notes provided by instructor.

Prerequisites: Basic knowledge in linear algebra, probability; Proficiency in Python coding

Grading

- Assignments: 15% (best 1 out of 2)
- Mini ML application project: 30%
- Mid-Term Exam: 25%,
- End-Term Exam: 30%

DS 285o Tensor Computations for Data Science (3:1)

Course Instructor

Prof. Ratikanta Behera, CDS

Course Description

This course is an introduction to tensor computations, focusing on theory, algorithms, and applications of tensor decompositions to data sciences. In the era of Big Data, Artificial Intelligence, and Machine Learning, we are faced with the need to process multiway (tensor-shaped) data. These data are mostly in the three or higher order dimensions, whose order of magnitude can reach billions. Huge volumes of multi-dimensional data are a great challenge for processing and analyzing; the matrix representation of data analysis is not enough to represent all the information content of the multiway data in different fields. Further, the importance of being able to decompose a tensor is (at least) two-fold. First, finding the decomposition provides hidden information about the data at hand, and second, having a concise decomposition of the tensor allows us to store it much more efficiently.

The course will provide an understanding of tensor operations and decomposition with various applications, including image deblurring, image compression, neural network, and solving high dimensional partial differential equations.

Learning Outcomes

- Students will learn basic understanding of the theoretical foundations of tensors computation.
- Students will choose efficient tensor decomposition for solving a specific problem.
- Students will learn efficient algorithms for tensor operations; including multiplication, decomposition, and inverse of tensors.
- Students will learn how to implement and use tensor computation in data sciences; including image deblurring, image compression, and solving high dimensional partial differential equations.
- Students will learn the basic tensor computation in the neural networks.

Syllabus

Unit-1 Fundamentals: Basic concepts of matrix properties: norms, rank, trace, inner products, Kronecker product, similarity matrix. Fast Fourier transform, diagonalization of matrices. Toeplitz and circulant matrices with their properties (eigenvalue and eigenvector), block matrix computation, and warm-up algorithms.

Unit-2 Introduction to Tensors: Tensors and tensor operations: Mode-n product of a tensor. Kronecker product of two tensors, tensor element product, Khatri-Rao product, the outer product. The Einstein product and t-product tensors. The explicit examples include identity tensor, symmetric tensor, orthogonal tensor, tensor rank, and block tensor.

Unit-3 Tensor Decomposition: Block tensor decomposition, Canonical Polyadic (CP) decomposition, the Tucker decomposition, the multilinear singular value (the higher-order SVD or HOSVD) decomposition, the hierarchical Tucker (HT) decomposition, and the tensor-train (TT) decomposition. Eigenvalue decomposition and singular value decomposition via t-product and the Einstein product. Truncated tensor singular value decomposition. Tensor inversion, and Moore-Penrose inverse. power tensor, solving system of multilinear equations.

Unit-4 Applications of Tensor Decompositions: Low-rank tensor approximation, background removal with robust principal tensor component analysis, image deblurring, image compression, compressed sensing with robust Regression, higher-order statistical moments for anomaly detection, solving elliptic partial differential equations.

Unit-5 Tensors for Deep Neural Networks: Deep neural networks, Tensor networks and their decompositions, including, CP decomposition, Tucker decomposition, Hierarchical Tucker decomposition, Tensor train and tensor ring decomposition, Transform-based tensor decomposition. Compressing deep neural networks.

Textbooks / References

- Books
 - Tensors for Data Processing: Theory, Methods, and Applications by Liu, Y. (Ed.). Academic Press. (2021)
 - Tensor Computation for Data Analysis by Liu Y, Liu J, Long Z, Zhu C. Springer, 2022
- Recent Articles
 - Block Tensor Unfoldings by S. Ragnarsson and C. F. Van Loan. *SIAM J. Matrix Anal. Appl.*, 33(1):149–169, 2012
 - Orthogonal Tensor Decompositions by T. G. Kolda. *SIAM Journal on Matrix Analysis and Applications*, 23(1):243–255, 2001
 - Tensor Decompositions and Applications by T. G. Kolda and B. W. Bader. *SIAM Rev.*, 51(3):455–500, 2009
 - An Order-p tensor Factorization with Applications in Imaging. *SIAM J Sci Comput.* 2013;35(1): A474–90
 - Solving multilinear systems via tensor inversion by M. Brazell, N. Li, C. Navasca, et al. *SIAM J. Matrix Anal Appl.* 2013;34(2):542–570

Prerequisites

Basic linear algebra with basic programming skills (in any programming language)

Grading

- Homework and Quizzes: 50%
- Project: 30%
- End-Term Exam: 20%

E0 295o Introduction to Cryptography (3:1)

Course Instructors

Prof. Bhavana Kanukurthi, CSA

Prof. Chaya Ganesh, CSA

Course Description

An introductory course in cryptography which will equip the students with the underlying mathematical foundations of the subject ranging from encryption to zero-knowledge.

Syllabus

Modular arithmetic and basic NT, Information theory, one time pad, Shannon impossibility, Computational relaxations, negligible functions, proofs by reduction, PRG intro, PRG defn, next-bit unpred and indistinguishability, equivalence of defns, PRF defn, PRG to PRF GGM, CPA, Cyclic groups, subgroups, structure of \mathbb{Z}_N , \mathbb{Z}_p^* , Assumptions: dlog, DDH, CDH, PRF candidate, Integrity, MACs, CCA defn, construction, Schnorr, Elgamal, RSA, signatures (OTS, FDH), Random oracle model, commitments, Identification schemes, sigma protocols, zero-knowledge proofs

Textbooks / References

1. Introduction to Modern Cryptography by Jonathan Katz and Yehuda Lindell [[Link](#)]
2. A Graduate Course in Applied Cryptography by Dan Boneh and Victor Shoup [[Link](#)]
3. Foundations of Cryptography by Oded Goldreich [[Link](#)]

Prerequisites: Discrete Mathematics, Basic Probability

Grading

HW + Mid-Term Exam + End-Term Exam + Class Participation

E1 252o Linear and Non-Linear Optimization (3:1)

Course Instructor

Prof. Chandramani Singh, ESE

Course Description

In this course, we will study the basics of linear and nonlinear optimization. We will also see several usages of optimization techniques in supervised and unsupervised learning.

Syllabus

Optimization examples, The basics - global vs local optimality; Convex sets, Convex and concave functions; First-order and second-order optimality conditions; Gradient descent methods, Conjugate gradient method, Newton method, Gradient projection method; Constrained optimization with equality and inequality constraints, Duality; Linear programming, simplex method, duality; Barrier and penalty function methods; Sub gradient descent methods; Proximal gradient descent; Augmented Lagrangian methods

Textbooks / References

1. Nonlinear Programming by D. Bertsekas, Athena Scientific, 2016
2. Linear and Nonlinear Programming by D. Luenberger and Y. Ye, Springer, 2008
3. Convex Optimization by S. Boyd and L. Vandenberghe, Cambridge University Press, 2004

Prerequisites: None

Grading

- Homework (assigned approximately once in two weeks): 40%
- Mid-Term Exam: 30%
- End-Term Exam: 30%

Homeworks will be assigned approximately once in two weeks, and a scanned copy of the solutions need to be turned in.

E1 285o: Advanced Deep Representation Learning (3:1)

Course Instructor

Prof. Prathosh A P, ECE

Course Description

This course contains discussions on cutting-edge topics on Advanced Deep learning. Given their practical applicability, it is imperative that the algorithms discussed during the lectures are duly implemented (as assignments) by the students. For compute, students can utilize freely-available resources such as Google Colab or Kaggle on toy/small-scale datasets to understand the implementation nuances of different algorithms. There will be 2 assignments and one project component along with terms papers that count for the laboratory credit.

Syllabus

Recap on Fundamentals of Deep Learning: Empirical Risk Minimization, Divergence minimizations and Likelihood maximization Techniques, Deep Learning Architectures (Convolutional and Recurrent Architectures).

Deep Generative Models: Introduction to Generative models, Autoregressive and invertible models, Latent variable models, Variational inference, and recognition networks (VAE, WAE), Adversarial Learning, Generative Adversarial networks, and variants (BiGAN, CycleGAN, StyleGAN, WGAN), Normalizing Flows, Score/Diffusion based models

Transfer Learning and Domain Adaptation: Discrepancy-Based Approaches: statistical (MMD) geometrical and architectural criteria, Generative Domain Adaptation: Adversarial and Non-adversarial Methods, Reconstruction based methods, Domain Generalization: Representation, data manipulation and Learning strategy methods

Few-shot and Meta Learning: Introduction to Multi-task and Transfer learning, Meta-learning framework for few-shot learning, Metric learning, comparators and relational networks, Optimization-based meta learning, Generative meta learning

Semi and Self-supervised Learning: Consistency Regularization, Proxy-label Methods, Active Learning, Weakly supervised learning methods, Self-supervised and Contrastive Representation Learning, Contrastive losses, Memory-bank techniques, BYOL, SWAV, SimCLR, MoCo, Hard negative mining.

Textbooks / References

1. Understanding Machine Learning: From Theory to Algorithms by Shai Ben-David and Shai Shalev-Shwartz, Cambridge University Press
2. Probabilistic Machine Learning: Advanced Topics by Kevin P. Murphy, MIT Press, 2023
3. Deep Learning by Aaron Courville, Ian Goodfellow, and Yoshua Bengio, MIT Press, 2016
4. Mathematics for Machine Learning by Deisenroth, Marc Peter, A. Aldo Faisal, and Cheng Soon Ong, Cambridge University Press, 2020
5. Machine Learning from Weak Supervision: An Empirical Risk Minimization Approach by Masashi Sugiyama, Han Bao, Takashi Ishida, Nan Lu, Tomoya Sakai, and Gang Niu, MIT Press
6. Deep Generative Modeling by Jakub M. Tomczak, Springer 2022
7. Semi-Supervised Learning by Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien, MIT Press
8. Seminal and Survey Papers from Machine Learning Conferences such as ICML, Neurips, ICLR, CVPR, AISTATS etc.

Prerequisites

1. *Mandatory:* A course on probability theory
2. *Mandatory:* A course on classical machine learning fundamentals
3. *Mandatory:* Moderate programming skills in Python

Grading

- 1 Minor and 1 Major: 50%
- Project and Assignments: (20+20) %
- Term Paper: 10%

E1 2860: Deep Generative Models (3:1)

Course Instructor

Prof. Prathosh A P, ECE

Course Description

This graduate-level course provides an in-depth exploration of deep generative models, including their probabilistic foundations and learning algorithms. Students will learn about various types of deep generative models such as variational autoencoders, generative adversarial networks, autoregressive models, normalizing flow models, energy-based models, and score-based models. The course will cover both theoretical foundations and practical implementations of these models using popular frameworks like PyTorch. Students will gain hands-on experience through lectures and assignments, allowing them to explore deep generative models across various AI tasks.

Course Objectives

- Develop a deep understanding of the importance of generative models in artificial intelligence and machine learning
- Design and implement generative models using popular frameworks
- Implement a range of generative models, including autoregressive models, normalizing flow models, energy-based models, and score-based models
- Build problem-solving skills by tackling challenges and complexities in the practical implementation of generative models

Core Competencies

- Deep Generative Models
- Autoregressive Models
- Variational Autoencoders
- Generative Adversarial Networks (GANs)
- Energy-Based Models
- Score-Based Models
- Diffusion Models
- Large Language Models

Syllabus

Recap on Fundamentals of Deep Learning:

Empirical Risk Minimization, Divergence Minimizations and Likelihood Maximization Techniques, Deep Learning Architectures (Convolutional, Recurrent and Transformer Architectures)

Deep Generative Models:

Introduction to Generative Models, Latent Variable Models, Variational Inference and Recognition Networks (Variational Auto Encoders and Wasserstein's Autoencoders)

Adversarial Learning, Generative Adversarial Networks, and variants (BiGAN, CycleGAN, StyleGAN, WGAN)

Diffusion and Score Based Models (DDPM, Consistency Models, DALLE, Stable Diffusion)

Auto-regressive and masked Language Models:

Large Language Models (Encoder/Decoder Models, BERT, GPT and other family)

Multi-Modal Models:

CLIP, Vision-Language Generative Models

Adaptation and Transfer:

Scaling LLMs (Mixture of Experts), Fine Tuning and Adaptation Techniques (LoRA, RAG)

Domain Adversarial Networks, Image-to-Image Translation, and other applications

Textbooks / References

1. Understanding Machine Learning: From Theory to Algorithms by Shai Ben-David and Shai Shalev-Shwartz, Cambridge University Press
2. Probabilistic Machine Learning: Advanced Topics by Kevin P. Murphy, MIT Press, 2023
3. Deep Learning by Aaron Courville, Ian Goodfellow, and Yoshua Bengio, MIT Press, 2016
4. Mathematics for Machine Learning by Deisenroth, Marc Peter, A. Aldo Faisal, and Cheng Soon Ong, Cambridge University Press, 2020
5. Machine Learning from Weak Supervision: An Empirical Risk Minimization Approach by Masashi Sugiyama, Han Bao, Takashi Ishida, Nan Lu, Tomoya Sakai, and Gang Niu, MIT Press
6. Deep Generative Modeling by Jakub M. Tomczak, Springer 2022
7. Deep Learning: Foundations and Concepts, by Christopher Bishop and Hugh Bishop, Springer, 2023
8. Seminal and Survey Papers from Machine Learning Conferences such as ICML, Neurips, ICLR, CVPR, AISTATS etc.

Prerequisites

1. Proficiency in Python: Coding assignments will be in Python.
2. Calculus and Linear Algebra: Students should be comfortable taking (multivariable) derivatives and understand matrix/vector notation and operations.
3. Probability Theory: Students should be familiar with basic probability distributions (Continuous, Gaussian, Bernoulli, etc.) functions.

Grading

- 2 Exams: 20% each (40)
- 3 Programming Assignments: 15% each (45)
- 3 Quizzes: 5% each (15)

E1 294o Edge and Cloud Systems for Machine Learning Algorithms (3:1)

Course Instructor

Prof. Sumit Kumar Mandal, CSA

Course Description

Increasing complexity of machine learning (ML) algorithms have necessitated the emergence of specialized computer systems. Some ML algorithms are executed at edge, some are executed on the cloud. In this course we will delve into different computing kernels of ML algorithms and see how they can be efficiently computed.

Syllabus

- Introduction to Deep Neural Networks (DNN), Convolutional Neural Networks (CNN)
- DNN processing on GPUs
- GPU input pipeline for DNN
- Accelerators for DNNs
- DNN accelerators with emerging technologies
- Accelerators for Graph Convolutional Networks (GCN)
- GCN accelerators with emerging technology
- Accelerators for Large Language Models (LLMs)

The course will involve moderate to heavy programming. The programming part will include building DNNs from scratch, understanding various performance analysis tools in Linux, training ML models with and without GPUs, parallel programming in Python. The programming assignments will be key to understanding ML models from systems perspective.

Textbooks / References

Efficient Processing of Deep Neural Networks by Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, Joel Emer, Morgan & Claypool Publishers

Prerequisites: No prerequisites as such

Grading

- Class Participation: 10%
- Scribe: 10%
- Quiz: 20%
- Homework: 20%
- Project: 40%

E1 3160 Deep Learning for Robotics (2:1)

Course Instructors

Dr. Shishir N. Y. Kolathaya, Dr. Pushpak Jagtap, Dr. Ravi Prakash, RBCCPS

Course Description

This graduate course will explore the new area of interaction between learning and control specifically applied to robotic systems, both from a foundational level together with a view toward application. The course will first build the necessary framework in which to understand basics of robotic systems, including robot kinematics and dynamics, sensing and estimation, control, machine learning, reinforcement learning and imitation learning. With these fundamentals the course will focus on data driven approaches for learning and control applied to robotic systems.

Syllabus

Robot dynamics and kinematics, nonlinear control and stability, Markov decision processes, reinforcement learning, imitation learning, model free methods, exploration vs. exploitation

Note: Active student participation will be expected in course projects, presentations, and delivering lectures in some of the selected topics.

Textbooks / References

1. A Mathematical Introduction to Robotic Manipulation by Richard M. Murray, Zexiang Li, S. Shankar Sastry, CRC Press, 1994
2. Reinforcement Learning: An Introduction by Richard S. Sutton and Andrew G. Barto, MIT Press, 2017
3. Deep Reinforcement Learning by Sergey Levine: <http://rail.eecs.berkeley.edu/deeprlcourse/>
4. Robot Modeling and Control by Mark W. Spong, Seth Hutchinson, M. Vidyasagar, Wiley, 2005

Prerequisites: Students must be well versed with basic mathematical concepts like linear algebra, classical analysis, and probability theory.

Grading

- Homeworks/Quiz: 60%
- Project Reports: 10%
- Project Presentations: 30%

E2 203o Wireless Communications (3:1)

Course Instructor

Prof. Sudhan Majhi, ECE

Course Description

Follow this link to know more: <https://ece.iisc.ac.in/~nextgenwrl/>

Syllabus

- Wireless channel modeling
- Diversity techniques to combat fading
- OFDM
- Capacity of wireless channels
- MIMO: Channel modeling, and transmit and receive architectures
- 4G/5G cellular system design
- Link budget analysis

Textbooks / References

1. Fundamentals of Wireless Communication by David Tse, Pramod Viswanath, Cambridge University Press, 2005
2. Wireless Communications by Andrea Goldsmith, Cambridge University Press, 2005

Prerequisites: Digital Communications

Grading

- Homeworks: 30% (Assigned about once in two weeks. Scanned pages of the solutions need to be submitted by students for evaluation)
- Mid-Term Exam: 30%
- End-Term Exam: 40%

E2 251o Communication Systems Design (3:1)

Course Instructor

Prof. A. Chockalingam, ECE

Course Description

The course gives a comprehensive treatment of how to design communication systems. Emphasis will be given to communication link design principles and communication transceiver design fundamentals. The course will introduce and elaborate on transmit/receive architectures and subsystem requirements/specifications for modern communication transceivers. A key component in the course is the modeling of impairments in the transmit and receive chains. Algorithms for the compensation of these impairments using digital processing techniques in the baseband is another highlight in the course. The role of deep neural networks (DNNs) in communication systems design will be introduced. Design of visible light communication links using LEDs and photodiodes will also be covered. There will be MATLAB programming exercises and assignments.

Syllabus

- Communication link design for AWGN channels; path loss models, noise figure, receiver sensitivity, pink noise, link budget for deep space communication - a case study
- Communication subsystem requirements/specifications and impairments modeling: analog front-end, up/down conversion architectures, oscillator phase noise, sampling jitter, IQ imbalance, carrier frequency offset (CFO), sampling frequency offset (SFO), error vector magnitude (EVM), DAC/ADC interface, quantization noise and clipping, dynamic range, ADC noise floor, non-linearities: 1 dB compression point, IIPn, power amplifier (PA) non-linearities, AM/AM, AM/PM
- Compensation techniques: Tx/Rx IQ imbalance estimation and compensation, effect of CFO in OFDM, CFO estimation and compensation in OFDM, SFO estimation and correction in OFDM
- Visible light wireless communications (VLC): transmitter, channel, receiver, VLC link design, MIMO-VLC

Textbooks / References

1. RF Analog Impairments Modeling for Communication Systems Simulation: Application to OFDM-based Transceivers by Lydi Smaini, John-Wiley & Sons, 2012
2. Wireless Receiver Architectures and Design: Antenna, RF, Synthesizers, Mixed Signal and Digital Signal Processing by Tony J. Roupheal, Academic Press, 2014
3. RF Transceiver Design for MIMO Wireless Communications by Abbas Mohammadi and Fadhel M. Ghannouchi, Springer-Verlag, 2012
4. Research papers in journals and conferences

Prerequisites: None

Grading

- MATLAB programming assignments (4 assignments, 10 marks each): 40%
- Mid-Term Exam: 30%
- End-Term Exam: 30%

E2 2760 Process Technology and System Engineering (3:1)

Course Instructor

Prof. Hardik Pandya, ESE

Course Description

In many engineering and science areas, miniaturization of the devices presents several challenges, and it requires expertise from several fields, giving rise to interdisciplinary research. A unique opportunity for manufacturing devices within the nm–mm range is possible due to the recent innovations in microfabrication. The six-orders-of-magnitude dimensional span fabricates several devices, including MOSFETs, microfluidics, microheaters, and other microscale sensors, with applications ranging from electronics to computation, clinical research, and robotics. Understanding this research area will also help students think and understand how to design novel devices based on the applications and potentially develop innovative solutions for several existing problems. With a knowledge of microfabrication and systems engineering to integrate the sensors within a data acquisition platform, a student/researcher can work in several fields, from microelectronics, robotics, electronics systems engineering, and instrumentation engineering shows the importance of this course. This course offers exposure to the right mix of concepts in microfabrication for fabricating several microscale electronic devices and their interfacing with data acquisition systems. The course comprehensively explores this exciting field of microengineering technologies, enabling you to understand, design, and build cutting-edge sensor solutions. The course is designed to equip students with the necessary knowledge and skills to fabricate, analyze, and integrate sensors effectively, setting a path to excel in multiple industries.

Syllabus

Introduction and Overview of Microfabrication Process Technology, Overview of Processes used in the Fabrication of Microsensors and Devices, Silicon Wafers - From Sand to the Laboratory, Thin Film Growth and Deposition Techniques, Fundamentals of Material Deposition Techniques: Overview of Physical Vapor Deposition (Sputtering, Thermal and E-beam evaporation) and Chemical Vapor Deposition (LPCVD, PECVD, MOCVD, ALD); Doping - Diffusion and Ion implantation techniques, Lithography – Modes of Lithography, Photoresists: Positive and Negative Photoresists, Masks: Bright Field and Dark Field Masks, Tools in Photolithography, Etching – Wet and Dry Etching, Fabrication of Micro-Engineered Devices: Process Flow of Device Fabrication using Semiconductor Wafers, Data Acquisition Systems Integrated with Signal Conditioning Circuits for Interfacing Sensors and Devices, Case Studies (sensors and transducers)

Lab Component (Videos): Familiarization with Gowning Procedure and Safety Protocols, Introductory Clean Room Visit and Overview of Equipment, Wafer Cleaning Processes, Thermal Evaporation of Metals, E-beam Evaporation of Metals and Insulators, Photolithography, Wet Etching of Metals, Semiconductors and Insulators, Device Fabrication: From Si to Microchips

Textbooks / References

1. Fundamentals of Microfabrication by Marc J Madou
2. Silicon VLSI Technology: Fundamentals, Practice and Modeling by James D. Plummer, Michael D. Deal, Peter D. Griffin
3. Fundamentals of Semiconductor Fabrication by Gary S. May and Simon M. Sze
4. VLSI Fabrication Principles: Silicon and Gallium Arsenide by Sorab K. Ghandhi
5. VLSI Technology by S. M. Sze
6. Fundamentals of Microelectronics by Behzad Razavi
7. Design with Operational Amplifiers and Analog Integrated Circuits by Sergio Franco, 2002, New York: McGraw-Hill
8. Sensors and Signal Conditioning by Ramon Pallás-Areny and John G. Webster, 2012, John Wiley & Sons

Course Outcomes

- Ability to understand the microfabrication process
- Understand the fabrication of PMOS and NMOS transistors
- Understand the fabrication of micro-cantilevers, flexible force sensors, inter-digitated electrodes, and microfluidic platforms
- Design and develop implantable devices for brain-computer interface
- Understand the design criteria and challenges in interfacing micro engineered devices with the real-world using data acquisition systems
- Learn several processes used in microfabrication through experiments in class 1000 and class 10000 cleanroom
- Designing Electronic Modules and acquiring skills in 3D printing

Prerequisites: Undergraduate in Electronics/Physics/Instrumentation/Mechanical

Grading

- Assignments: 20%
- Mid-Term Exam: 30%
- Final-Term Exam: 50%

E2 287o Communication Networking Lab (1:1)

Course Instructor

Prof. Parimal Parag, ECE

Syllabus

- End-to-end traffic flow
 - Understanding the data flow in a typical web application – Introduction to Wireshark
- Introduction to network configuration
 - Creating topology in mininet
 - Performance measurements such as throughput and delay
- Performance of transport protocols
 - Performance measurement for UDP/TCP in present of packet drops, delay, and contention
 - Introduction to CUBIC, BBR, DCTCP
- Introduction to Python
 - Control statements, loops, data structures
 - File IO, Network IO
- Introduction to binary classification
 - SVM and Kernel methods
 - k nearest neighbours
- Introduction to network traffic
 - Network traffic capture
 - Traffic classification

E3 280o Semiconductor Devices for Nanoelectronics (3:1)

Course Instructor

Prof. Kausik Majumdar, ECE

Syllabus

Module 1: Introduction

Module 2: Basics of quantum mechanics

Module 3: Electrons in solids

Module 4: Crystal lattice, band structure

Module 5: Lattice vibrations and phonons

Module 6: Doping, junctions and related devices, carrier transport

Module 7: MOSFET

Module 8: Device fabrication steps in nanoelectronics

Textbooks / References

1. Introduction of Quantum Mechanics by D. J. Griffiths, Prentice Hall
2. Quantum Mechanics by A. Ghatak and S. Lokanathan, Trinity Press
3. Quantum Mechanics by V. K. Thankappan, New Age
4. Solid State Physics by N. W. Ashcroft and N. D. Mermin
5. Physics of Semiconductor devices by S. M. Sze, Wiley-Interscience
6. Fundamentals of modern VLSI devices by Y. Taur and T. H. Ning, Cambridge University Press

Prerequisites: None

Grading

- Assignments: 30%
- Mid-Term Exam: 20%
- End-Term Exam: 30%
- Mini Project: 20%

E8 204o Antenna Theory and Practice (3:1)

Course Instructor

Prof. Debdeep Sarkar, ECE

Course Description

Antennas are essential components for the wireless communication and sensing systems, which require transition from guided EM (electromagnetic) waves to radiated EM waves and vice versa. MIMO (multiple-input multiple-output) antenna technology is one of the key enablers for high data-rate communication systems (5G and beyond). Alongside MIMO antennas deployed for handsets, base stations, and wireless access points, full-duplex (FD) radios have the potential of achieving increased spectral efficiency and reduced latency, which is essential for futuristic 6G systems. Besides usefulness in the context of communication systems, MIMO and FD concepts have vital importance in present-day Radar technology, especially automotive Radars that are essential for ADAS (advanced driver assistance systems). In this course, the emphasis will be put on the design and signal processing aspects involving MIMO and FD antenna systems for both communication and Radar technology.

Syllabus

- **Antenna Fundamentals:** Understanding EM radiation, Wave-equation and solution, Hertzian dipoles, Brief Overview of Dipoles, Monopoles, Slot and Microstrip patch antennas, understanding basics of reflection coefficients, radiation patterns, and polarization, Friis equation and path-loss model
- **Understanding Antenna Arrays:** Array factor, Steering vector, Uniform and non-uniform arrays, Difference between MIMO and phased arrays
- **Design of MIMO Antennas for Handsets and Base-Stations:** Spatial, pattern and polarization diversity, calculation of ECC, TARC, CCL
- **Design of Antennas for Full-Duplex Communication:** Overview of various SIC (self-interference cancellation) schemes, the importance of passive cancellation, mutual coupling fundamentals, different mutual coupling reduction techniques for FD antennas
- **MIMO Radars:** Comparison with Phased Array Radars, Antenna Design, Series-fed antennas, Signal processing and design using MATLAB

The course will have programming and design assignments using MATLAB Antenna Toolbox for understanding and visualization.

Textbooks / References

1. Field and Wave Electromagnetics by D. K. Cheng, Pearson Education Asia Ltd, Second Edition, 2006
2. Antenna Theory - Analysis and design by C. A. Balanis, John Wiley, Fourth Edition, 2016
3. Antenna Theory and Design by W. L. Stutzman and G. A. Thiele, John Wiley & Sons Inc, 1981
4. Antennas by J. D. Karus, McGraw Hill, 1988
5. Microstrip antennas by I. J. Bahl and P. Bhartia, Artech house, 1980

Prerequisites

Presence of preliminary knowledge about vectors, coordinate transform, partial differential equations, circuit theory and transmission lines will be great. However, most of these topics will be reviewed before introducing any new topic in the class.

Grading

- Home Assignments: 30%
- Mid-Term Exam: 20%
- End-Term Exam/Course Project: 50%

E8 242o Radio Frequency Integrated Circuits and Systems (3:1)

Course Instructor

Prof. K. J. Vinoy, ECE

Course Description

This is a course on high frequency circuits and systems with applications to wireless systems.

Syllabus

Fundamentals: Review of Transmission line Theory, Impedance transformation, Terminated transmission lines; Smith chart, impedance matching; Board level Planar (Microstrip and Coplanar waveguide) implementations, Introduction to Maxwell Eqn, Wave eqn, Wave propagation; types; Introduction to Waveguides, Modes of propagation, substrate integrated waveguides.

Microwave network analysis, ABCD-, S-, X-parameters; Conformal analysis of planar transmission lines; Even/Odd Mode analysis for Passive RF components and circuits;

Passive Components: Power dividers; Couplers; Resonators at RF/microwaves; Image parameter and insertion loss methods, Transmission line implementations; Low pass filters, Bandpass and band stop filter; Theory of Periodic Structures, Microwave filters with Periodic structures; Microwave metamaterials, CRLH Transmission lines; Applications

Antennas: Introduction, definitions and basic principles of Microwave antennas, Design of typical antennas for Wireless systems; Antenna applications of metamaterials; Antennas arrays; Design of planar antennas

Active Circuits: Diode based RF circuits: diplexers, varactors; biasing; Circuits for RF energy harvesting; Basics of high frequency amplifier design, biasing techniques, simultaneous tuning of 2 port circuits; Noise in RF Circuits; designs based on impedance match noise performance; LNA Design; Analysis of non-linearities in RF devices; Amplifier Design using S-parameters of devices: LNA Design, Power amplifier design

Textbooks / References

1. Microwave Engineering by D.M. Pozar, John Wiley
2. Microwave and RF Wireless Systems by DM Pozar, Artech House
3. Foundations for Microwave Engineering by R.E. Collin, John Wiley
4. Elements of Electromagnetics by MNO Sadiku, 3rd Ed., Oxford Univ. Press
5. Antenna Theory: Analysis and Design by C.A. Balanis, John Wiley
6. Journal/Magazine Articles

Prerequisites: None

Grading

- Multiple Quiz: 40%
- Design and Analysis exercises using commercial tools: 30%
- Final Project: 30%

E9 246o Advanced Topics in Image Processing (3:1)

Course Instructor

Prof. Soma Biswas, EE

Course Description

This course will serve as a smooth transition from the traditional image processing/computer vision techniques to the state-of-the-art deep learning-based techniques. At the end of this course, the students will have a broader perspective of the field and the different real-world applications like object detection, domain adaptation, etc. will help the students understand how to use the techniques taught for real-world applications. The assignments and projects will give them hands-on experience. There will be 3 Lecture hours per week, and coding assignments and projects are a part of this course.

Syllabus

- **Image Features:** Local features, Corner detectors, Scale-Invariant Feature Transform (SIFT) and its variants, Shape detectors like Hough Transform, Texture descriptors like LBP, Evaluation
- **Deep Learning of Image Features:** Convolution layers, fully connected layers, back propagation, non-linear activation, regularization, batch normalization, drop-out, network architectures
- **Image Segmentation:** Graph-based techniques, Active Shape Models, Active Appearance Models, deep learning-based segmentation, evaluation
- **Object Detection:** RCNN and its improved variants like Fast-RCNN, Faster RCNN, YOLO, Evaluation
- **Self-Supervised Learning:** Contrastive Learning, RotNet, etc.
- **Selected Topics:** Vision Transformers for image classification, Multi-modal foundation model like CLIP, few real-world applications, generative models

Textbooks / References

7. Computer Vision: A Modern Approach by David A. Forsyth & Jean Ponce, Prentice Hall
8. Deep Learning by Ian Goodfellow, Yoshua Bengio and Aaron Courville, MIT Press, 2016
9. Recent papers, which will be provided when the topics are taught in class.

Prerequisites: Digital Image Processing (E9 241o)

Grading

- Assignments: 35%
- Mid-Term Exam: 15%
- End-Term Exam: 25%
- Final Project: 25%