Master of Science in ECE - Machine Learning & Data Science Focus

COURSE SCHEDULE

Core Coursework (16 units)

- ECE269: Linear Algebra
- ECE271A: Statistical Learning I
- ECE 289: Probability and Statistics for Data Science
- ECE188: Programming for Data Analysis

Select four courses (16 units) from the following areas. At least one course per area.

Analytics:

- ECE271B: Statistical Learning II
- ECE273: Convex Optimization and Applications
- ECE275A: Parameter Estimation I

Computation:

- ECE289: Optimization and Acceleration of Deep Learning on Various Hardware Platforms
- ECE289: Scalable Learning
- ECE289: Software for Data Science
- ECE289: Parallel Processing in Data Science

Applications:

- ECE204: Statistical Learning in Neuroscience
- ECE207: Computational Evolutionary Biology
- ECE276A: Sensing & Estimation in Robotics
- ECE276B: Planning & Learning in Robotics
- ECE276C: Advances in Robot Manipulation
- ECE285: Machine Learning for Physical Applications
- ECE267: Security of Hardware Embedded System
- ECE289: Big Network Data

Technical Electives (16 units)

- Any 4 unit, 200+ course from ECE, CSE, MAE, BENG, CENG, NANO, SE, MATS MATH, PHYS or CogSci taken for a letter grade may be counted. Exceptions to this list require departmental approval.
- Up to 12 units of undergraduate ECE coursework (ECE 111+ only) may be counted
- M.S. Students (Plan II) are allowed no more than 4 units of 299 as technical electives. Ph.D. and M.S. Students (Plan I) are allowed no more than 8 units of 299 as technical electives.

COURSE	Title	Required Elective	Fall	Winter	Spring	Notes
			REQUIRED			
ECE269	Linear Algebra	Required	Pal	Kim		
ECE271A	Statistical Learning I	Required	Vasconcelos			
ECE289	Probability and Statistics for Data Science	Required	Orlitsky	.		
ECE180	Programming for Data Analysis	Required	Staff	Staff	Staff	
			ANALYTICS (At least one course)			
ECE271B	Statistical Learning II	Elective		M. Vasconcelos		Prerequisite: 271A
ECE273	Convex Optimization and Applications	Elective			Pal	
ECE275A	Parameter Estimation I	Elective	Kreutz-Delgado			
			COMPUTATION (At least one course)			
ECE289	Optimization and Acceleration of Deep Learning on Various Hardware Platforms	Elective		Farinaz		
ECE289	Scalable Learning	Elective			Staff	
ECE289	Software for Data Science	Elective		Staff		
ECE289	Parallel Processing in Data Science	Elective			An	

APPLICATIONS (At least one course)						
ECE204	Statistical Learning for Biosignal Processing	Elective				
ECE207	Computational Evolutionary Biology	Elective			Mirarab	
ECE276A	Sensing & Estimation in Robotics	Elective	Atanasov			
ECE276B	Planning & Learning in Robotics	Elective		Atanasov		Prerequisite: ECE276A
ECE276C	Advances in Robot Manipulation	Elective			Yip	Prerequisitie: ECE276A
ECE285	Machine Learning for Physical Applications	Elective			Gerstoft	
ECE268	Security of Hardware Embedded System	Elective		Koushanfar		
ECE289	Big Network Data	Elective	Franceschetti			

Machine Learning - Data Science Course Information

Fundamentals (4 required courses)

Linear algebra - ECE-269

Syllabus: Linear algebra (finite dimensional) is the study of mathematical principles guiding matrix operations with particular reference to the geometry of vector spaces over which such operations are defined. In this course, we will build the mathematical foundations of linear algebraic techniques which will justify their use in diverse applications in signal processing, communication, machine learning and network sciences. Topics include fundamentals of vector and Hilbert spaces (bases, subspace, inner product, norm, orthogonality), projection and least squares, systems of linear equations (overdetermined vs. underdetermined) and role of sparsity, eigenanalysis and linear dynamical systems, Hermitian matrices and variational characterization, singular value decomposition (SVD) and principal component analysis, positive semidefinite matrices.

Machine learning ECE-271A
Syllabus: Bayesian decision theory; parameter estimation; maximum likelihood; the bias-variance trade-off; Bayesian estimation; the predictive distribution; conjugate and noninformative priors; dimensionality and dimensionality reduction; principal component analysis; Fisher's linear discriminant analysis; density estimation; parametric vs. kernel-based methods; expectation-maximization; applications.

Probability and statistics (somewhere in between 109 and 250)
Syllabus: Probabilistic models, random variables, common distributions, expectations, Markov chains, random walks, law of large numbers, central limit theorem, maximum likelihood, inference, confidence intervals, hypothesis testing, linear regression.

Programming for data analysis
Syllabus: A hands-on course designed to teach students Python and its usage in Data Science applications. Topics include:

- Understand Python object-oriented and functional programming styles
- Learn key scientific computing packages
- Apply key Python data structures and algorithms effectively
- Enhance productivity with Python development workflows
- Develop deployable codes using modern package management and source control

Advanced (At least 4 courses, at least one from each category)

Analytics

Advanced machine learning ECE-271B
Syllabus: Linear discriminants; the Perceptron; the margin and large margin classifiers; learning theory; empirical vs. structural risk minimization; the VC dimension; kernel functions; reproducing kernel Hilbert spaces; regularization theory; Lagrangian optimization; duality theory; the support vector machine; boosting; Gaussian processes; applications.

Convex Optimization and Applications ECE-273
Syllabus: This course will introduce the mathematical principles guiding modern convex optimization methods. These include the geometry of convex sets, behavior of convex functions, existence and characterization of optimal solutions of convex problems via duality theory. It will focus on recognizing and formulating convex problems, model/relax a seemingly "non-convex" problem in terms of a tractable convex problem, and understand the properties of optimal solution(s) using duality. It covers applications in a variety of fields (system design, signal processing, machine learning and pattern recognition, combinatorial optimization, financial engineering, etc.).

Parameter estimation ECE-275A
Syllabus: Linear least Squares (batch, recursive, total, sparse, pseudo-inverse, QR, SVD); Statistical figures of merit (bias, consistency, Cramer-Rao lower-bound, efficiency); Maximum likelihood estimation (MLE); Sufficient statistics; Algorithms for computing the MLE including the Expectation Maximization (EM) algorithm. The problem of missing information; the problem of outliers. The Bayesian statistical framework; Parameter and state estimation of Hidden Markov Models, including Kalman Filtering and the Viterbi and Baum-Welch algorithms.

Computation

Optimization and Acceleration of Deep Learning on Various Hardware Platforms
Syllabus: This course focuses on a holistic end-to-end methodology for optimizing the physical performance metrics of Deep Learning on hardware platforms, e.g., real-time performance, energy, memory, and power. The hardware platforms include CPU-CPU, CPU-GPU, and CPU-FPGA architectures. We start by discussing the hardware characteristics and the effect of the architecture on the DL performance. We will cover platform-specific algorithm and data transformation that contribute to significant improvement in deep learning performance.

Scalable learning
Syllabus: Scalable methods for running algorithms on distributed systems, including Hadoop and Spark.

Parallel Processing in Data Science
Syllabus: High performance computing, parallel programming, graphical processing units (GPU's), CUDA language and libraries, with application in Data Science.

Software for Data Science
Syllabus: TBD

Applications

Statistical Learning for Biosignal Processing ECE-204
Syllabus: Medical device systems increasingly measure biosignals from multiple sensors, requiring computational analyses of complex multivariate time-varying data. The combination of statistics and algorithms produces statistical learning methods that automate the analysis of complex data. Such machine learning methods are used to analyze data collected by medical devices to enhance their design as well as to customize their operation for individual patients. Example applications within the domain of neural engineering that utilized unsupervised and supervised generative statistical modeling techniques are explored. This course assumes familiarity with key statistical methods.
Prerequisites: ECE 271A or ECE 271B; graduate standing.

Computational evolutionary biology ECE-207
Syllabus: A hands-on course where students learn to apply a set of computational techniques to a real biological question, namely evolutionary biology (e.g., the study of tree-of-life). The course involves building

biological tools in assignments and projects and we focus on scalability to big genomic data. Techniques taught include dynamic programming, continuous time Markov models, hidden Markov models, statistical inference of phylogenies, sequence alignment, uncertainty (e.g., bootstrapping), heterogeneity (e.g., phylogenetic mixture models). Programming skills required.

Sensing & Estimation in Robotics ECE-276A
Syllabus: This course covers the mathematical fundamentals of Bayesian filtering and their application to sensing and estimation in mobile robotics. Topics include maximum likelihood estimation (MLE), expectation maximization (EM), Gaussian and particle filters, simultaneous localization and mapping (SLAM), visual features and optical flow, and hidden Markov models (HMM).
Prerequisites: equivalent of ECE101, 153, 171, 174

Planning & Learning in Robotics ECE- 276B
Syllabus: This course covers optimal control and reinforcement learning fundamentals and their application to planning and decision making in mobile robotics. Topics include Markov decision processes (MDP), Pontryagin's Maximum Principle, linear quadratic regulation (LQR), deterministic planning (A^* and RRT^*), value and policy iteration, Q-learning, and policy gradient methods
Prerequisite: ECE276A

Advances in Robot manipulation ECE-276C
Syllabus**:** Robot Manipulation involves the use of robot effectors (like arms, trunks, hands, etc.) to operate in real environments. It ranges from low-level control (such as how a robot should move its joints to move its gripper towards an object), to high-level decision making (such as whether the robot should make the move in the first place). Many useful algorithms that have been developed in the areas of control theory, artificial intelligence, and now machine learning are being used in unison to achieve tasks. This class is set up in a way to explore reinforcement learning as a means to solve challenging robot manipulation problems. . Part 1 will cover topics pertinent to robot manipulation and will rapidly focus on examining new algorithms for achieving more complex robot motions and behaviors. Part 2 will involve a substantial project component involving developing a new machine learning algorithm to solve some open challenges in robot manipulation.
Prerequisite: ECE 276A

Machine Learning for Physical Applications ECE-285
Syllabus: Machine learning has received enormous interest recently. However, for physical problems there is reluctance to use machine learning. Machine learning cannot replace existing physical models, but improve certain aspects of them. To learn from data, we use probability theory, which has been the mainstay of statistics and engineering for centuries. Probability theory can be applied to any problem involving uncertainty. The class will focus on implementations.
Offered: Spring

Security of Hardware Embedded System ECE-268
The course gives an overview of areas of security and protection of modern hardware, embedded systems, and IoTs. Covers essential cryptographic methodologies and blocks required for building a secure system. Topics include low overhead security, physical and side-channel attacks, physical security primitives, physical security and proofs of presence, hardware-based secure program execution, scalable

implementation of secure functions, emerging technologies, and rising threats. Recommended preparation: Programming in a standard programming language. Undergraduate level knowledge of the IC design flow and digital designs.

Big Network Data
People, societies, biological micro-organisms and man-made devices connect to each other and form all kinds of complex networks. Thanks to technological advancements, an ocean of data has become available describing these connections. How do we analyze these “Big Network Data” and construct relevant engineering models?

Network science is a new discipline that addresses this question, investigating the topology and dynamics of complex networks arising from massive data collection, and aims at explaining and predicting the emerging trends and features of real systems. These systems are modeled as a statistical ensemble of interacting components, capable of exhibiting emerging complexity as a network property.

The course focuses on both rigorous foundations as well as on getting practical hands-on experience in analyzing real-world network data leading to learning and prediction in a variety of domains, including social, economic, medical, and engineering domains.

Specific topics include network structure (percolation graphs, paths, diameter, chemical distance, small worlds); processes on network (interacting particle systems, community detection, segregation, contagion), statistical methods: (sampling, bayesian inference, learning, and intervention); constrained optimization (network formation and evolution)