

NJIT@JerseyCity - Spring 2021 Courses

Math 661 - Applied Statistics (Andrew Pole)

Extracting information from data is the application of statistical description, analysis, and inference. Businesses, industry, academia, governments face these questions every day, looking for opportunities and seeking solutions for emergent problems: How much bandwidth can that streaming compression technology free up? Will we gain customers if we improve the audio quality? How should we examine performance data from four different designs of a prosthetic limb? What is the reliability of different component manufacturing processes? How do we allocate funding to maximize relief impact? In this course you will learn basic and advanced mathematical and statistical tools and techniques for extracting information from your data, big or small, and to make quantitatively justified statements about general populations of concern (customers, patients, supply chain choke points, ...). You will also use real-world data from many settings in the current leading data analytic environments R and Python to learn how to address these business, research, and social needs.

IS 601 - Web Systems Development (Ryan Tolboom)

This course trains in the development of web-based systems using Python and the Django web framework. You will learn to program with Python and Javascript, currently two of the three [most in-demand programming languages](#), and learn to develop a web based system through an intensive, hands-on project that requires application of real-world problem-solving skills. You will also learn to utilize the Django web framework, one of the [ten most popular web frameworks](#) and currently in use by [Instagram and Pinterest](#). Model-View-Template design, the Django templating language, and object relational mapping for database access will be covered in this course. At the conclusion of the course, you should be able to create a bespoke web interface for a standard internal business application.

IS 650 - Data Visualization and Interpretation (Aritra Dasgupta)

"We're not just fighting an epidemic; we're fighting an infodemic" - The WHO has observed that we are currently facing unprecedented challenges due to data-driven misinformation about the ongoing pandemic which can have severe consequences. The key contributors to this "infodemic" are often poorly designed visual charts, which can disinform or misinform, and do so, at scale. The need to learn the basics of visualization design and interpretation, both as a data consumer and a data practitioner, is more relevant and consequential now than ever before. Visualization techniques, when designed effectively, have emerged as a game-changer in data science, business Intelligence, and journalism domains for communicating data-driven insights. In this course you will learn the theory and application of visual design and how to reason about what makes a visualization good or bad. You will learn how to employ principles imbibed from computer science, human perception, cognition and information design to make your charts error-free, effective, and actionable. You will learn why visualization is needed in the data science pipeline, how it can be used to solve real-world problems, and how to distinguish between "good", "bad", or "misleading" visualizations. You will gain hands-on experience in building interactive visual interfaces and dashboards using Python, Tableau and Javascript.

CS 506 - Foundations of Computer Science (Jing Li)

This course provides the mathematical and analytical foundations of computer science and its applications to various areas in CS. The course covers the material traditionally known as “discrete mathematics”, with special emphasis on CS applications and analysis of algorithms. Discrete Mathematics is viewed as the pillar or language of CS. One needs to be fluent in it to work in many CS fields including data science, machine learning, and software engineering. The course topics include sets and logic, proof techniques, proof by induction, functions and relations, analysis of algorithms, recursion, recurrence equations, divide-and-conquer design technique, counting methods (permutations and combinations), basic discrete probability, and if time permits, introduction to number theory and a brief introduction to graphs and trees.

CS 630 - Operating Systems Design (Jing Li)

Operating systems are the foundation of computer science and computer systems. A deep understanding of OS allows one to develop programs and applications that are efficient and scalable when running on the underlying hardware. This course prepares you with the knowledge of operating systems concepts and designs that appear in computer systems, via discussing the basic operating system abstractions, mechanisms, and their implementations. In particular, topics we will cover include operating system designs, multiprogramming, timesharing, resource scheduling, synchronization of concurrent processes, management of hierarchical storage, embedded and real-time systems, virtualization, and distributed systems.

CS 634 - Data Mining (Pantelis Monogioudis)

The process of extracting signals from massive amounts of noisy data is a multi-step endeavor that involves, data analysis, visualization, learning algorithms, model analysis and verification. In this course you will learn not only what methods to use at each one of these stages but also, how to design and architect end-to-end pipelines that deliver on your technical and, ultimately, business objectives. Using real-world use cases from Zillow, Uber and other technology companies, you will use Python and cutting-edge open source machine learning libraries and tools to mine massive, publicly available repositories of structured (tabular) as well as unstructured (e.g. natural language) data. The skills acquired in this course are essential for aspiring data analysts or scientists across all data-driven industries.

CS 636 - Data Analytics with R Programming (Yao Shen)

This course will teach how to program in R and how to use R for effective data analysis. Students will learn how to install and configure R necessary for an analytics programming environment and gain basic analytic skills via this high-level analytical language. The course covers fundamental knowledge in R programming. Popular R packages for data science will be introduced as working examples. It will familiarize you with the commonly used analytical techniques in data science and develop the way of data science thinking: learn how to preprocess, explore and interpret real data and learn how to model real problems using computational techniques. The format of the course will include lectures by the instructor, computing labs, class discussion, directed reading and student presentation or project.

CS 656 - Internet and Higher Layer Programming (Kumar Mani)

This course introduces the protocols and standards of the TCP/IP suite that govern the functioning of the Internet. The material covered in class is a top-down approach on introduction, discussion, and analysis of protocols from the data-link layer to the application layer. Alternative protocols to the TCP/IP suite and new protocols adopted by this suite are discussed. Numerical examples related to network planning and protocol functioning are analyzed.

CS 677 - Deep Learning (Ioannis Koutis)

Deep Learning (DL) is a subfield of Machine Learning that has delivered disruptive technologies, created AI algorithms that outperform humans in various tasks, and paves the way for broader advances in Science. DL consists of a set of specialized techniques that exploit the abundant availability of data and computational power to build models that are composed of multiple processing layers and learn representations of data at multiple levels of abstraction. Only a few years back, the development of DL models required significant expertise, but the introduction of open source DL libraries like TensorFlow and PyTorch has opened the area to scientists and professionals with more diverse backgrounds. The course opens with a review of Artificial Neural Networks that guides you through TensorFlow (or PyTorch) and enables you to build novel ANN architectures. Then it presents the evolution of progressively deeper architectures for Convolutional Neural Networks, that addressed various training difficulties and led to very successful image classification models. In this spatial context, you will also learn about Generative Adversarial Networks that are behind the fascinating 'Deep Fake' images and videos. The course then takes you to the emerging applications of Recurrent Neural Networks in temporal data, including Natural Language Processing. You will also learn how Deep Learning combined forces with Reinforcement Learning to create models that can play games. The course will also touch upon selected topics like the ability of deep networks to generalize, techniques for 'pruning' deep networks in order to make them more computationally efficient, and successful applications of DL methods in the Sciences.