NJIT@JerseyCity - Fall 2020 Courses

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 657   Spatio-Temporal Urban Analytics (Aritra Dasgupta)

This is a *project-based course* in which we will focus on visualization of high-dimensional data collected about urban populations. Students will have an option of working on an individual project or a group project where they will focus on visual communication of data-driven insights. To get the most out of this course, it is *necessary* that students have taken the *IS 650* course and are familiar with basic data visualization techniques.

Students will get to choose from a number of project options and the project goal can be tailored based on a match between their preferences and the course objectives. They will work closely with the instructor and his team of PhD students and the schedule (of meetings/deliverables) can be mutually agreed upon before the start of the course.

Students will be evaluated based on the progress they make towards meeting the project deliverables. The focus will not solely be on a final product but on the lessons they learn while committing and rectifying mistakes along the way.

The project options include, but are not limited to, the following topics: *data discovery*, where the goal is to semantically integrate openly available spatio-temporal data (e.g., NYC open data) for driving decision-making about various facets, like health, education, economy, etc., of an urban population; *educational data mining*, where openly available college and university datasets (across states and multiple years) will be combined with students' economic and demographic data for identifying potential bias and lack fairness in admissions' criteria; studying *information consumption*, where social media data (e.g., Twitter) combined with sentiment analysis can be used to analyze the nature of content that propagates through social networks in times of a global pandemic/"infodemic" and reflect on the need for each person be their own moderator for preserving their mental health.

All of the aforementioned projects comprise an element of research. They require students to exercise their critical thinking skills and will have ample opportunity for them to come up with innovative approaches/solutions, and in turn, contribute directly to the research projects.

IS 665   Data Analysis for Information Systems (Lin Lin)

This graduate level course introduces students to the world of data analytics from an information systems perspective, focusing on the application of various data analysis techniques in business practices. We cover a wide spectrum of topics ranging from fundamental statistics to database, data warehouse, data visualization, and data mining. Being an introductory course, our approach is "shallow and wide", emphasizing on giving students a complete view of the data analytics profession, covering as many different sub-areas as time allows while not diving too deep into any one specific domain. The goal is to serve as a "guided tour" for students to gain knowledge about the different sub-areas of data analytics and understanding of which area is a best fit for their personal developments. More in-depth materials and discussion for each sub-area will be provided upon students' requests. Course topics include the rudiments of probability and random variables, data visualization, data warehousing and OLAP analysis, dashboard, scorecard, data mining algorithms, optimization techniques, DSS and knowledge systems.

CS 631   Data Visualization and Interpretation (Ananya Dass)

A relational database organizes data into tables which can be linked—or related—based on data common between them, and you can communicate with these databases using the standard Structured Query Language (SQL). This course is about designing, creating and querying relational databases. As data becomes more and more important, we see an explosion of specialized databases (OLAP, time-series, document, graph, etc.), data processing tools (Hadoop, Spark, Flink), data buses (Kafka, RabbitMQ), etc. We also have more applications that rely on this data infrastructure, whether data visualization tools (Tableau, Grafana, PowerBI, Superset), web frameworks (Rails, Django) or custom-built data-driven applications, built on a complex stack, with infrastructure on the bottom and applications on top. To make this work, we need an interface that allows the different layers of this stack to communicate with one another. This can be done using SQL, and the course will teach you this technology and allow you to gain hands-on experience with SQL to design, create, and manipulate a relational database by developing a real-world application.

CS 644   Introduction to Big Data   (David Bader)

Data science is a rapidly emerging discipline at the intersection of machine learning, algorithms, and statistics. Solving real-world problems in domains such as financial services, cybersecurity, and health informatics requires the generation, storage, management, and analytics, on massive data sets. This course will take you on a journey of developing high performance data analytics for massive data sets using accelerated computer architectures, state-of-the-art software frameworks, and real-world case studies. In-depth topics will focus on mathematical modeling, scalable algorithms, and high-performance computing, for fast and practical solutions. Together we will explore real-world data modeling and prediction from the Data Science Design Manual and learn How 45 Successful Companies Used Big Data Analytics to Deliver Extraordinary Results.

CS 645   Security and Privacy in Computer Systems (Reza Curtmola)

Security has become pervasive in our daily use of computers and networks because of the increasing reliance of critical industries on computer systems and networks, the ever-growing software complexity and attack sophistication, and the need to ensure data security and privacy both in transit and when being stored. The course covers fundamental principles of building secure systems and techniques to ensure data security and privacy. Topics include authentication and access control mechanisms, physical security, operating and file system security, malicious code threats and software security, web security, database security, and recent security incidents. You will be exposed to the theoretical underpinnings of cybersecurity and also get a chance to apply these concepts in several projects.

CS 675 Intro to Machine Learning (Ioannis Koutis)

Machine Learning develops computer programs that can improve their performance by tapping into existing data and taking feedback from the environment. Systems based on ML have already exceeded human performance in several tasks, including image medical image classification and games like Chess and Go. ML has also made leaps in even more complicated tasks, like Natural Language Processing or self-driving vehicles, and it has even produced art that imitates the style of human artists! This course offers an intense introduction to the fundamental ML concepts and algorithms that constitute the core of these spectacular developments. It takes you on a tour from the basic mathematical notions and algorithms to some of the recent developments, e.g. Deep Networks or Recurrent Networks. You will gain exposure to cutting-edge ML development tools such as Scikit-learn and TensorFlow via hands-on assignments and projects that instill a working and immediately applicable knowledge of ML methods and will prepare you for more advanced ML courses.

CS 677 Deep Learning (Pantelis Monogioudis)

Deep learning for the curious mind, is a course that explores the methods and fascinating applications of deep neural networks. We build the subject from the ground up starting from the foundations of inductive & transductive learning and explain why stochastic gradient descent and back propagation (BP) were critical to successfully train deep neural networks. We will use tools such as regularization and hyperparameter optimization to help us train networks that have the best inferential performance when they face big numbers of high dimensional inputs. We will then learn two architectures that embed innate priors: Convolutional Neural Networks (CNN) spatially and Recurrent Neural Networks (RNN) temporally. They are key in enabling applications in computer vision and natural language processing where learning is all about building the right representations. This is often achieved using autoencoders that allow us to learn the useful properties of our data and generative modeling that allow us to generate data from the underlying target function. Equipped with these principles, we then start to combine deep learning with others disciplines to appreciate how they have disrupted three key application domains. (1) Combining deep learning and control lead to Deep Reinforcement Learning (DRL) and today's disruption in industrial automation via robotic manufacturing. (2) Combining deep learning and linguistics lead to breakthroughs in Natural Language Understanding (NLU) and our ability to retrieve information with the right context. (3) Combining deep learning with computer vision leads to disruption in transportation enabling self-driving cars and trucks to safely drive themselves in a wide range of challenging environments. Drawing from these three application areas we strike a good balance between concepts and hands-on experiments and assign projects that are challenging and at the same time a lot of fun!