A Focus in Artificial Intelligence
Course Catalog

Bridge Courses: Students without formal preparation in Computer Science, may count up to two bridge courses towards the focus.

CSE 581: Computer Science Fundamentals: Theory
The course consists of two parts. The first part covers discrete mathematics -- a division of mathematics that is extensively used in computer science. The topics covered include: logic (propositional logic and predicate logic), proof techniques, sequences (mathematical induction and recursion), and functions. The second part covers the theory of computation -- a division of theoretical computer science that deals with what can be computed and what cannot be computed on a computer. The topics covered include: computational models (FA, PDA, and Turing machines), grammars accepted by different computational models (regular grammars, context-free grammars, and unrestricted grammars), languages accepted by different computational models (regular languages, context-free language, and Turing-acceptable languages), Turing-complete systems, and algorithmically unsolvable problems. *3 credits, Letter graded (A, A-, B+, etc.)*

CSE 582: Computer Science Fundamentals: Data Structures and Algorithms
The course consists of two parts. The first part covers data structures to efficiently store, organize, modify, and access data. Topics include: arrays, stacks, queues, linked lists, trees, sets, hash maps, priority queues, and graphs. The second part covers the design and analysis of algorithms for solving computer science problems. Topics include: algorithm analysis, exhaustive search algorithms, divide and-conquer algorithms, greedy algorithms, and dynamic programming algorithms. *3 credits, Letter graded (A, A-, B+, etc.)*

CSE 583: Computer Science Fundamentals: Programming Abstractions
TBD

Required Courses: Students will choose at least one.

CSE 512: Machine Learning
Machine Learning is centered around automated methods that improve their own performance through learning patterns in data, and then use the uncovered patterns to predict the future and make decisions. Examples include document/image/handwriting classification, spam filtering, face/speech recognition, medical decision making, robot navigation, to name a few. See this for an extended introduction. This course covers the theory and practical algorithms for machine learning from a variety of perspectives. The topics include Bayesian networks, decision tree learning, Support Vector Machines, statistical learning methods and unsupervised learning, as well as theoretical concepts such as the PAC learning framework, margin-based learning, and VC dimension. Short programming assignments include hands-on experiments with various learning algorithms, and a larger course project gives students a chance to dig into an area of their choice. See the syllabus for more. This course is designed to give a graduate-level student a thorough grounding in the methodologies, technologies, mathematics, and algorithms currently needed by people who do research in machine learning. *3 credits, Letter graded (A, A-, B+, etc.)*

CSE 537: Artificial Intelligence
A comprehensive introduction to the problems of artificial intelligence and techniques for attacking them.

Topics include problem representation, problem-solving methods, search, pattern recognition, natural language processing, learning, expert systems, and AI programming languages and techniques. Covers both theoretical methods and practical implementations. *3 credits, Letter graded (A, A-, B+, etc.)*

Electives: Students will choose up to 3.

CSE 505: Computing with Logic
The course explores logic-based computing and logic programming. It includes an introduction to programming in logic, covering basic techniques for solving problems in a logic programming system. Particular attention will be paid to user interface issues and how a logic system can provide a useful computing environment. The course covers implementation issues, emphasizing how a logic programming system generalizes both traditional programming language systems and traditional database systems. *3 credits, Letter graded (A, A-, B+, etc.)*

CSE 519: Data Science Fundamentals
This course will cover the building blocks of data science from managing the data itself to algorithmic and analytical techniques. Specific topics include data preparation, exploratory data analysis, statistics, visualization, optimization, unstructured data, distributed analyses. *3 credits, Letter graded (A, A-, B+, etc.)*

CSE 525: Introduction to Robotics
This course introduces the fundamental concepts in robotics, including coordinate transformations, visual perception, sensors, path planning, kinematics, feedback control, and feedforward control. These topics will be exemplified with several state-of-the-art robotics platforms. The course will also focus on applying the fundamental concepts to the key approaches to mobile robot control (reactive, behavior-based, and hybrid), and briefly discuss robot learning and multi-robot systems. *3 credits, Letter graded (A, A-, B+, etc.)*

CSE 527: Computer Vision
Introduction to basic concepts in computer vision. Low-level image analysis, image formation, edge detection, segmentation. Image transformations for image synthesis methods for 3D scene reconstruction, motion analysis, object recognition. *3 credits, Letter graded (A, A-, B+, etc.)*

CSE 538: Natural Language Processing
The course offers an introduction to Natural Language Processing techniques and applications. The course introduces basic deep learning techniques for constructing representations of natural language texts, syntactic analyses, and canonical applications such as Question Answering, and Machine Translation. The emphasis is on understanding how (i) the basic requirements of language analyses and specific application needs can be formulated as learning problems, (ii) how aspects of natural language, such as syntactic and semantic structure, inspire the design of and are captured by modern state of the art techniques, and (iii) how to evaluate and understand the limitations and shortcomings of the models. *3 credits, Letter graded (A, A-, B+, etc.)*

CSE 544: Probability and Statistics for Data Scientists
The course will cover core concepts of probability theory and an assortment of standard statistical techniques. Specific topics will include random variables and distributions, quantitative research methods (correlation and regression), and modern techniques of optimization and matching learning (clustering and prediction). *3 credits, Letter graded (A, A-, B+, etc.)*

CSE 545: Big Data Analytics
The course will cover concepts and standard tools used to analyze, so called, Big Data. Specifically, it will cover algorithmic approaches to analyzing large datasets: MapReduce, graph analytics, text analytics, steaming algorithms, as well as modern distribution analysis platforms (e.g. Hadoop, Spark). *3 credits, Letter graded (A, A-, B+, etc.)*

CSE 564: Visualization
Visualization plays an increasingly important role in the understanding of the massive data that are nowadays being collected in almost any domain – science, medicine, business, commerce, finance, social networks, and many more. As such, visualization is often deeply integrated into the analytics tools developed for data science. This course will discuss both foundations and applications of this emerging paradigm known as visual analytics. It will begin with the basics – visual perception, cognition, human-computer interaction, the sense-making process, data mining, computer graphics, and information visualization/It will then move to discuss how these elementary constituents are coupled into an effective visual analytics pipeline that allows humans to interactively reason with data and gain insight. *3 credits, Letter graded (A, A-, B+, etc.*