

EN.601 (COMPUTER SCIENCE)

Courses

EN.601.104. Computer Ethics. 1 Credit.

Students will examine a variety of topics regarding policy, legal, and moral issues related to the computer science profession itself and to the proliferation of computers in all aspects of society, especially in the era of the Internet. The course will cover various general issues related to ethical frameworks and apply those frameworks more specifically to the use of computers and the Internet. The topics will include privacy issues, computer crime, intellectual property law – specifically copyright and patent issues, globalization, and ethical responsibilities for computer science professionals. Work in the course will consist of weekly assignments on one or more of the readings and a final paper on a topic chosen by the student and approved by the instructor. The ePortfolio tag(s) on this course signify that there are one or more assignments offered in the course that provide students with the opportunity to be assessed for proficiency in completion of the relevant ePortfolio requirement(s).

Distribution Area: Humanities

EN Foundational Abilities: Writing ePortfolio (FA.1.1eP), Ethical Reflection (FA5)

EN.601.124. The Ethics of Artificial Intelligence and Automation. 3 Credits.

The expansion of artificial intelligence (AI)-enabled use cases across a broad spectrum of domains has underscored the benefits and risks of AI. This course will address the various ethical considerations engineers need to engage with to build responsible and trustworthy AI-enabled autonomous systems. Topics to be covered include: values-based decision making, ethically aligned design, cultural diversity, safety, bias, AI explainability, privacy, AI regulation, the ethics of synthetic life, and the future of work. Case studies will be utilized to illustrate real-world applications. Students will apply learned material to a group research project on a topic of their choice. The ePortfolio tag(s) on this course signify that there are one or more assignments offered in the course that provide students with the opportunity to be assessed for proficiency in completion of the relevant ePortfolio requirement(s).

Distribution Area: Humanities, Engineering

AS Foundational Abilities: Science and Data (FA2)

EN Foundational Abilities: Ethical Reflection (FA5), Ethical Reflection ePortfolio (FA5eP)

EN.601.164. Human and Machine Intelligence Alignment. 3 Credits.

The challenge of ensuring that the actions of individuals and systems – whether human or machine – are consistent with shared goals, reflect our values, and promote societal well-being is known as “the alignment problem.” Over millennia, humans have developed many coordination and cooperation “technologies” – such as customs, values, norms, laws, organizations, governments, and markets—that partially solve the problem of human intelligence alignment. As we develop and deploy advanced technologies like artificial intelligence, we are similarly concerned that their use is consistent with shared goals, reflect our values, and promote societal well-being. In this course we will explore the parallels between human intelligence alignment and machine intelligence alignment to help engineers and technologists become reflective practitioners who can grapple wisely with the alignment problem broadly understood. The ePortfolio tag(s) on this course signify that there are one or more assignments offered in the course that provide students with the opportunity to be assessed for proficiency in completion of the relevant ePortfolio requirement(s).

Distribution Area: Humanities, Engineering

EN Foundational Abilities: Writing ePortfolio (FA.1.1eP), Ethical Reflection (FA5), Ethical Reflection ePortfolio (FA5eP)

EN.601.220. Intermediate Programming. 4 Credits.

This course teaches intermediate to advanced programming, using C and C++. (Prior knowledge of these languages is not expected.) We will cover low-level programming techniques, as well as object-oriented class design, and the use of class libraries. Specific topics include pointers, dynamic memory allocation, polymorphism, overloading, inheritance, templates, collections, exceptions, and others as time permits. Students are expected to learn syntax and some language specific features independently. Course work involves significant programming projects in both languages.

Prerequisite(s): (A grade of C+ is required in EN.500.112 OR EN.500.113 OR EN.500.114) OR (AP Computer Science OR EN.500.132 OR EN.500.133 OR EN.500.134) Students will not be able to register until their grades are posted for the prerequisite courses.

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.226. Data Structures. 4 Credits.

This course covers the design, implementation and efficiencies of data structures and associated algorithms, including arrays, stacks, queues, linked lists, binary trees, heaps, balanced trees and graphs. Other topics include sorting, hashing, Java generics, and unit testing. Course work involves both written homework and Java programming assignments.

Prerequisite(s): A grade of C+ or better in EN.500.112 OR EN.601.220 OR AP Computer Science OR EN.500.132. Students can't register until grades for prerequisites are posted.

Distribution Area: Engineering, Quantitative and Mathematical Sciences

AS Foundational Abilities: Science and Data (FA2)

EN.601.229. Computer System Fundamentals. 3 Credits.

This course covers modern computer systems from a software perspective. Topics include binary data representation, machine arithmetic, assembly language, computer architecture, performance optimization, memory hierarchy and cache organization, virtual memory, Unix systems programming, network programming, and concurrency. Hardware and software interactions relevant to computer security are highlighted. Students will gain hands-on experience with these topics in a series of programming assignments.

Prerequisite(s): EN.601.220

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.230. Mathematical Foundations for Computer Science. 4 Credits.

This course provides an introduction to mathematical reasoning and discrete structures relevant to computer science. Topics include propositional and predicate logic, proof techniques including mathematical induction, sets, relations, functions, recurrences, counting techniques, simple computational models, asymptotic analysis, discrete probability, graphs, trees, and number theory.

Prerequisite(s): Student may not enroll if taken EN.601.231 OR EN.601.431;EN.500.112 OR EN.500.113 OR EN.500.114 OR EN.500.132 OR EN.500.133 OR EN.500.134 OR EN.601.220

Distribution Area: Engineering, Quantitative and Mathematical Sciences
AS Foundational Abilities: Science and Data (FA2)

EN.601.231. Automata & Computation Theory. 3 Credits.

This course is an introduction to the theory of computing. topics include design of finite state automata, pushdown automata, linear bounded automata, Turing machines and phrase structure grammars; correspondence between automata and grammars; computable functions, decidable and undecidable problems, P and NP problems, NP-completeness, and randomization.

Prerequisite(s): Students may not enroll if taken EN.601.230.;EN.550.171/EN.553.171 OR EN.553.172

Distribution Area: Engineering, Quantitative and Mathematical Sciences
AS Foundational Abilities: Science and Data (FA2)

EN.601.257. Computer Graphics and 3D Game Programming. 3 Credits.

In this course, students will program a game of their own design using an off-the-shelf game engine while learning about the 3D computer graphics concepts behind the engine's components. Classes will consist of a mix of theory and practice. The theory will be presented through lectures on topics including transformations, lighting, shading, shape representations, spatial querying and indexing, animation, and special effects. Practice will involve in-class programming exercises and contributions to the game project with periodic in-class presentations of progress to date. Students are expected to have a strong programming background and to be familiar with basic linear algebra concepts.

Prerequisite(s): EN.601.220 AND EN.601.226 AND (AS.110.201 OR AS.110.212 OR EN.553.291 OR EN.553.295)

Distribution Area: Engineering

EN.601.264. Practical Generative AI. 3 Credits.

This course is a comprehensive guide for students eager to explore the world of generative AI and its practical applications in software development. Designed with a hands-on approach, it equips you with the foundational knowledge of generative AI, introduces a suite of AI development tools, and covers key AI platforms and frameworks. You'll gain the skills needed to build and deploy AI-powered applications, culminating in a substantial team project that offers real-world experience in creating AI-driven software. By the end of the course, you'll be prepared to integrate AI into your applications and development process, unlocking new avenues for creativity and innovation.

Recommended course background: EN.601.290 or EN.601.490.

Prerequisite(s): Students who have previously taken, or are currently enrolled in, EN.601.470 OR EN.601.670 OR EN.601.471 OR EN.601.671 are not eligible to take EN.601.264.;EN.601.220 AND EN.601.226 AND EN.601.280

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.270. Open Source Software Engineering (Semesters of Code I). 3 Credits.

The course will provide students a development experience focused on learning software engineering skills to deliver software at scale to a broad community of users associated with open source licensed projects. The class work will introduce students to ideas behind open source software with structured modules on recognizing and building healthy project structure, intellectual property basics, community & project governance, social and ethical concerns, and software economics.

Prerequisite(s): EN.601.220 AND EN.601.226

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.277. Disinformation Self-Defense. 3 Credits.

Scientific, statistical and logical literacy is a necessary skill for evaluating policy proposals, reading news articles with an appropriately critical eye, and making informed choices as consumers and voters. Misunderstanding of claims made in scientific publications, online publishing platforms, and mass media drives, in part, the spread of malicious misinformation and propaganda online. Further, many actors have the means, the motive and the opportunity to mislead the public in a variety of subtle and not so subtle ways. This class will give you tools to discern valid and invalid forms of inference and discourse, and give you tools to communicate precisely, argue appropriately, and stay on top of research and news with an appropriately skeptical attitude. A use case used throughout the class will be online disinformation surrounding the COVID-19 pandemic. The class will draw on historical and modern literature on linguistic, logical, and probabilistic fallacies, statistical and logical inference, data visualization, cognitive biases, and the scientific method.

Prerequisite(s): EN.601.230 OR EN.553.171 OR EN.553.172 OR AS.150.118

Distribution Area: Quantitative and Mathematical Sciences, Social and Behavioral Sciences

AS Foundational Abilities: Science and Data (FA2)

EN.601.280. Full-Stack JavaScript. 3 Credits.

A full-stack JavaScript developer is a person who can build modern software applications using primarily the JavaScript programming language. Creating a modern software application involves integrating many technologies - from creating the user interface to saving information in a database and everything else in between and beyond. A full-stack developer is not an expert in everything. Rather, they are someone who is familiar with various (software application) frameworks and the ability to take a concept and turn it into a finished product. This course will teach you programming in JavaScript and introduce you to several JavaScript frameworks that would enable you to build modern web, cross-platform desktop, and native/hybrid mobile applications. A student who successfully completes this course will be on the expedited path to becoming a full-stack JavaScript developer.

Prerequisite(s): Students must not have taken or be concurrently enrolled in EN.601.421 or EN.601.621 Object Oriented Software Engineering.;EN.601.220 OR EN.601.226

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.290. User Interfaces and Mobile Applications. 3 Credits.

This course will provide students with a rich development experience, focused on the design and implementation of user interfaces and mobile applications. A brief overview of human computer interaction will provide context for designing, prototyping and evaluating user interfaces. Students will invent their own mobile applications and implement them using the Android SDK, which is JAVA based. An overview of the Android platform and available technologies will be provided, as well as XML for layouts, and general concepts for effective mobile development. Students will be expected to explore and experiment with outside resources in order to learn technical details independently. There will also be an emphasis on building teamwork skills, and on using modern development techniques and tools.

Prerequisite(s): EN.601.220 AND EN.601.226

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.315. Databases. 3 Credits.

Introduction to database management systems and database design, focusing on the relational and object-oriented data models, query languages and query optimization, transaction processing, parallel and distributed databases, recovery and security issues, commercial systems and case studies, heterogeneous and multimedia databases, and data mining.

Prerequisite(s): Students can only take one of the following: EN.601.315, EN.601.415, OR EN.601.615.;EN.601.226

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.318. Operating Systems. 3 Credits.

This course covers fundamental topics related to operating systems theory and practice. Topics include processor management, storage management, concurrency control, multi-programming and processing, device drivers, operating system components (e.g., file system, kernel), modeling and performance measurement, protection and security, and recent innovations in operating system structure. Course work includes the implementation of operating systems techniques and routines, and critical parts of a small but functional operating system.

Prerequisite(s): Students may receive credit for only one of the following: EN.601.318, EN.601.418, OR EN.601.618.;EN.601.226 AND EN.601.229

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.340. Web Security. 3 Credits.

This course begins with reviewing basic knowledge of the World Wide Web, and then exploring the central defense concepts behind Web security, such as same-origin policy, cross-origin resource sharing, and browser sandboxing. It will cover the most popular Web vulnerabilities, such as cross-site scripting (XSS) and SQL injection, as well as how to attack and penetrate software with such vulnerabilities. Students will learn how to detect, respond, and recover from security incidents. Newly proposed research techniques will also be discussed. Note: This undergrad version will not have the same paper component as the other versions of this course.

Prerequisite(s): Students may receive credit for only one of 601.340/440/640.;EN.601.226 AND EN.601.229 AND EN.601.280

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.350. Genomic Data Science. 3 Credits.

This course will use a project-based approach to introduce undergraduates to research in computational biology and genomics. During the semester, students will take a series of large data sets, all derived from recent research, and learn all the computational steps required to convert raw data into a polished analysis. Data challenges might include the DNA sequences from a bacterial genome project, the RNA sequences from an experiment to measure gene expression, the DNA from a human microbiome sequencing experiment, and others. Topics may vary from year to year. In addition to computational data analysis, students will learn to do critical reading of the scientific literature by reading high-profile research papers that generated groundbreaking or controversial results. Recommended Course Background: Knowledge of the Unix operating system and programming expertise in a language such as Perl or Python.

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.356. Seminar: Computer Integrated Surgery II. 1 Credit.

Students may receive credit for EN.601.456 or EN.601.356, but not both. Lecture only version of EN.601.456 (no project). Recommended Course Background: EN.601.455 or instructor permission required.

Prerequisite(s): EN.601.455 or instructor permission.;Students may receive credit for either EN.601.356 or EN.601.456, but not both.

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.402. Digital Health and Biomedical Informatics. 1 Credit.

Advances in technology are driving a change in medicine, from personalized medicine to population health. Computers and information technology will be critical to this transition. We shall discuss some of the coming changes in terms of computer technology, including computer-based patient records, clinical practice guidelines, and region-wide health information exchanges. We will discuss the underlying technologies driving these developments - databases and warehouses, controlled vocabularies, and decision support.

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.404. Brain & Computation. 1 Credit.

Computational and network aspects of the brain are explored. The topics covered include structure, operation and connectivity of neurons, general network structure of the neural system, and the connectivity constraints imposed by pre- and post-natal neural development and the desirability of network consistency within a species. Both discrete and continuous aspects of neural computation are covered. Precise mathematical tools and analyses such as logic design, transient and steady state behavior of linear systems, and time and connectivity randomization are discussed. The concepts are illustrated with several applications. Memory formation from the synaptic level to the high level constructs are explored. Students are not expected to master any of the mathematical techniques but are expected to develop a strong qualitative appreciation of their power. Cerebellum, which has a simple network connectivity, will be covered as a typical system.

Prerequisite(s): Students can receive credit for EN.601.404 or EN.601.604, but not both;(EN.553.291 OR ((AS.110.201 OR AS.110.212) AND AS.110.302)) AND (EN.553.420 OR EN.553.421 OR EN.553.211 OR EN.553.310 OR EN.553.311) AND EN.601.433

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.411. Computer Science Innovation & Entrepreneurship II. 3 Credits.

This course is the second half of a two-course sequence and is a continuation of course EN.660.410.01, CS Innovation and Entrepreneurship, offered by the Center for Leadership Education (CLE). In this sequel course the student groups, directed by CS faculty, will implement the business idea which was developed in the first course and will present the implementations and business plans to an outside panel made up of practitioners, industry representatives, and venture capitalists.

Prerequisite(s): EN.660.410

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.413. Software Defined Networks. 3 Credits.

Software-Defined Networks (SDN) enable programmability of data networks and hence rapid introduction of new services. They use software-based controllers to communicate with underlying hardware infrastructure and direct traffic on a network. This model differs from that of traditional networks, which use dedicated hardware devices (i.e., routers and switches) to control network traffic. This technology is becoming a key part of web scale networks (at companies like Google and Amazon) and 5G/6G networks. Its importance will keep on growing. Many of today's services and applications, especially when they involve the cloud, could not function without SDN. SDN allows data to move easily between distributed locations, which is critical for cloud applications. A major focus will be on how this technology will be used in 5G and 6G Networks. The course will cover basics of SDN, ongoing research in this area, and the industrial deployments.

Prerequisite(s): Students can receive credit for EN.601.413 or EN.601.613, but not both; EN.601.414 OR EN.601.614

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.414. Computer Networks. 3 Credits.

Topics covered will include application layer protocols (e.g. HTTP, FTP, SMTP), transport layer protocols (UDP, TCP), network layer protocols (e.g. IP, ICMP), link layer protocols (e.g. Ethernet) and wireless protocols (e.g. IEEE 802.11). The course will also cover routing protocols such as link state and distance vector, multicast routing, and path vector protocols (e.g. BGP). The class will examine security issues such as firewalls and denial of service attacks. We will also study DNS, NAT, Web caching and CDNs, peer to peer, and protocol tunneling. Finally, we will explore security protocols (e.g. TLS, SSH, IPsec), as well as some basic cryptography necessary to understand these. Grading will be based on hands-on programming assignments, homeworks and two exams.

Prerequisite(s): EN.601.226 AND EN.601.229 or permission.; Students may receive credit for only one of EN.600.344, EN.600.444, EN.601.414, EN.601.614.

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.415. Databases. 3 Credits.

Similar material as EN.601.315 covered in more depth for advanced undergraduates. Introduction to database management systems and database design, focusing on the relational and object-oriented data models, query languages and query optimization, transaction processing, parallel and distributed databases, recovery and security issues, commercial systems and case studies, heterogeneous and multimedia databases, and data mining. (www.cs.jhu.edu/~yarowsky/cs415.html)

Prerequisite(s): Students may receive credit for only one of the following: EN.601.315, EN.601.415, OR EN.601.615.; EN.601.226

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.418. Operating Systems. 3 Credits.

Similar material as EN.601.318, covered in more depth. Intended for advanced undergraduate students. This course covers fundamental topics related to operating systems theory and practice. Topics include processor management, storage management, concurrency control, multi-programming and processing, device drivers, operating system components (e.g., file system, kernel), modeling and performance measurement, protection and security, and recent innovations in operating system structure. Course work includes the implementation of operating systems techniques and routines, and critical parts of a small but functional operating system.

Prerequisite(s): Students may receive credit for only one of the following: EN.601.318, EN.601.418, OR EN.601.618.; EN.601.226 AND EN.601.229

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.419. Cloud Computing. 3 Credits.

Clouds host a wide range of the applications that we rely on today. In this course, we study common cloud applications, traffic patterns that they generate, critical networking infrastructures that support them, and core networking and distributed systems concepts, algorithms, and technologies used inside clouds. We will also study how today's application demand is influencing the network's design, explore current practice, and how we can build future's networked infrastructure to better enable both efficient transfer of big data and low-latency requirements of real-time applications. The format of this course will be a mix of lectures, discussions, assignments, and a project designed to help students practice and apply the theories and techniques covered in the course. Prerequisites: EN.601.226 or permission. Students can only receive credit for one of 601.419/619. Recommended: a course in operating systems, networks or systems programming.

Prerequisite(s): Students may earn credit for EN.601.419 or EN.601.619, but not both.; EN.601.226 (or EN.600.226) AND EN.601.414 or permission from the instructor.

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.420. Parallel Computing for Data Science. 3 Credits.

This course studies parallelism in data science, drawing examples from data analytics, statistical programming, and machine learning. It focuses mostly on the Python programming ecosystem but will use C/C++ to accelerate Python and Java to explore shared-memory threading. It explores parallelism at all levels, including instruction level parallelism (pipelining and vectorization), shared-memory multicore, and distributed computing. Concepts from computer architecture and operating systems will be developed in support of parallelism, including Moore's law, the memory hierarchy, caching, processes/threads, and concurrency control. The course will cover modern data-parallel programming frameworks, including Dask, Spark, Hadoop!, and Ray. The course will not cover GPU deep-learning frameworks nor CUDA. The course is suitable for second-year undergraduate CS majors and graduate students from other science and engineering disciplines that have prior programming experience and familiarity with Python.

Prerequisite(s): EN.601.226 AND EN.601.229; Students may receive credit for only one of EN.600.320, EN.600.420, EN.601.320, EN.601.420, EN.601.620.

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.421. Object Oriented Software Engineering. 3 Credits.

This course covers object-oriented software construction methodologies and their application. The main component of the course is a large team project on a topic of your choosing. Course topics covered include object-oriented analysis and design, UML, design patterns, refactoring, program testing, code repositories, team programming, and code reviews.

Prerequisite(s): Students may receive credit for only one of EN.600.321, EN.600.421, EN.601.421, EN.601.621.; EN.601.220 AND EN.601.226 AND (EN.601.280 OR EN.601.290)

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.422. Software Testing & Debugging. 3 Credits.

Studies show that testing can account for over 50% of software development costs. This course presents a comprehensive study of software testing, principles, methodologies, tools, and techniques. Topics include testing principles, coverage (graph coverage, logic coverage, input space partitioning, and syntax-based coverage), unit testing, higher-order testing (integration, system-level, acceptance), testing approaches (white-box, black-box, grey-box), regression testing, debugging, delta debugging, and several specific types of functional and non-functional testing as schedule/interest permits (GUI testing, usability testing, security testing, load/performance testing, A/B testing etc.). For practical topics, state-of-the-art tools/techniques will be studied and utilized.

Prerequisite(s): EN.601.290 OR EN.601.421; Students can take EN.601.422 or EN.601.622, but not both.

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.424. Reliable Software Systems. 3 Credits.

Reliability is an essential quality requirement for all artifacts operating in the real-world, ranging from bridges, cars to power grids. Software systems are no exception. In this computing age when software is transforming even traditional mission-critical artifacts, making sure the software we write is reliable becomes ever more important. This course exposes students to the principles and techniques in building reliable systems. We will study a set of systematic approaches to make software more robust. These include but are not limited to static analysis, testing framework, model checking, symbolic execution, fuzzing, and formal verification. In addition, we will cover the latest research in system reliability.

Prerequisite(s): EN.601.220 AND (EN.601.328 OR EN.601.428)

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.425. Software System Design. 3 Credits.

This course introduces modern software systems design, with an emphasis on how to design large-scale systems, assess common system design trade-offs, and tackle system design challenges. It covers non-functional requirements, API design, distributed systems concepts, modern software building blocks (e.g., load balancers, caches, containers, etc.). Additionally, it includes case studies of common system design problems, some drawn from interview questions. Ultimately, this course helps learners become better software engineers.

Prerequisite(s): Students may receive credit for EN.601.425 OR EN.601.625, but not both.; EN.601.315 OR EN.601.415 OR EN.601.615 OR EN.601.280 OR EN.601.290 OR EN.601.340 OR EN.601.440 OR EN.601.640 OR EN.601.421 OR EN.601.621

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.426. Principles of Programming Languages. 3 Credits.

Functional, object-oriented, and other language features are studied independent of a particular programming language. Students become familiar with these features by implementing them. Most of the implementations are in the form of small language interpreters. Some type checkers and a small compiler will also be written. The total amount of code written will not be overly large, as the emphasis is on concepts. The ML programming language is the implementation language used. Prerequisites include EN.601.226. No Freshmen or Sophomores.

Prerequisite(s): Students can receive credit for EN.601.426 or EN.601.626, but not both; EN.601.226

Distribution Area: Engineering, Quantitative and Mathematical Sciences

AS Foundational Abilities: Science and Data (FA2)

EN.601.428. Compilers & Interpreters. 3 Credits.

Introduction to compiler design, including lexical analysis, parsing, syntax-directed translation, symbol tables, run-time environments, and code generation and optimization. Students are required to write a compiler as a course project. Recommended background: EN.601.230 or EN.601.231.

Prerequisite(s): EN.601.226 AND EN.601.229

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.429. Functional Programming in Software Engineering. 3 Credits.

How can we effectively use functional programming techniques to build real-world software? This course will primarily focus on using the OCaml programming language for this purpose. Topics covered include OCaml basics, modules, standard libraries, testing, quickcheck, build tools, functional data structures and efficiency analysis, monads, streams, and promises. Students will practice what they learn in lecture via functional programming assignments and a final project.

Prerequisite(s): Students can receive credit for EN.601.429 or EN.601.629, but not both.;EN.601.226 OR Instructor Permission
Distribution Area: Engineering
AS Foundational Abilities: Science and Data (FA2)

EN.601.430. Combinatorics & Graph Theory in Computer Science. 3 Credits.

This course covers the applications of combinatorics and graph theory in computer science. We will start with some basic combinatorial techniques such as counting and pigeon hole principle, and then move to advanced techniques such as the probabilistic method, spectral graph theory and additive combinatorics. We shall see their applications in various areas in computer science, such as proving lower bounds in computational models, randomized algorithms, coding theory and pseudorandomness.

Prerequisite(s): EN.601.230 OR EN.553.171 OR EN.553.172 OR EN.550.171;Students may receive credit for EN.601.430 or EN.601.630, but not both.
Distribution Area: Engineering, Quantitative and Mathematical Sciences
AS Foundational Abilities: Science and Data (FA2)

EN.601.431. Theory of Computation. 3 Credits.

This course covers the theoretical foundations of computer science. Topics included will be models of computation from automata to Turing machines, computability, complexity theory, randomized algorithms, inapproximability, interactive proof systems and probabilistically checkable proofs. Students may not take both 601.231 and 601.431/601.631, unless one is for an undergrad degree and the other for grad.

Prerequisite(s): Students who have taken EN.601.631 OR EN.601.231 are not eligible to take EN.601.431.;EN.553.171 OR EN.553.172 OR EN.601.230
Distribution Area: Engineering, Quantitative and Mathematical Sciences
AS Foundational Abilities: Science and Data (FA2)

EN.601.433. Intro Algorithms. 3 Credits.

This course concentrates on the design of algorithms and the rigorous analysis of their efficiency. topics include the basic definitions of algorithmic complexity (worst case, average case); basic tools such as dynamic programming, sorting, searching, and selection; advanced data structures and their applications (such as union-find); graph algorithms and searching techniques such as minimum spanning trees, depth-first search, shortest paths, design of online algorithms and competitive analysis.

Prerequisite(s): EN.601.226 AND (EN.553.171 OR EN.553.172 OR EN.601.230 OR EN.601.231);Students may receive credit for only one of EN.600.363, EN.600.463, EN.601.433, EN.601.633.
Distribution Area: Engineering, Quantitative and Mathematical Sciences
AS Foundational Abilities: Science and Data (FA2)

EN.601.434. Randomized and Big Data Algorithms. 3 Credits.

This course will cover fundamental methods of randomization for efficient algorithms and relevant techniques in probabilistic analysis. The first part of the course will discuss classical randomized algorithms, including the randomized algorithm for the min-cut problem, hashing techniques, tail inequalities, and the probabilistic method. The second part will delve into advanced topics such as coresets methods for clustering algorithms, the Johnson-Lindenstrauss lemma, and the applications of streaming and sketching algorithms.

Prerequisite(s): Students may receive credit for only one of the following: EN.601.434 OR EN.601.634.;(EN.601.433 OR EN.601.633) AND (EN.553.211 OR EN.553.310 OR EN.553.311 OR EN.553.420 OR EN.553.421)

Distribution Area: Engineering, Quantitative and Mathematical Sciences
AS Foundational Abilities: Science and Data (FA2)

EN.601.435. Approximation Algorithms. 3 Credits.

This course provides an introduction to approximation algorithms. Topics include vertex cover, TSP, Steiner trees, cuts, greedy approach, linear and semi-definite programming, primal-dual method, and randomization. Additional topics will be covered as time permits. There will be a final project. Students may receive credit for EN.601.435 or EN.601.635, but not both.

Prerequisite(s): Students can only receive credit for one of EN.601.435 or EN.601.635.;EN.600.363 OR EN.601.433 OR EN.601.633 OR permission.
Distribution Area: Engineering, Quantitative and Mathematical Sciences
AS Foundational Abilities: Science and Data (FA2)

EN.601.436. Algorithmic Game Theory. 3 Credits.

This course provides an introduction to algorithmic game theory: the study of games from the perspective of algorithms and theoretical computer science. There will be a particular focus on games that arise naturally from economic interactions involving computer systems (such as economic interactions between large-scale networks, online advertising markets, etc.), but there will also be broad coverage of games and mechanisms of all sorts. Topics covered will include a) complexity of computing equilibria and algorithms for doing so, b) (in)efficiency of equilibria, and c) algorithmic mechanism design.

Prerequisite(s): EN.600.363 OR EN.600.463 OR EN.601.433 OR EN.601.633
Distribution Area: Engineering, Quantitative and Mathematical Sciences
AS Foundational Abilities: Science and Data (FA2)

EN.601.438. Algorithmic Foundations of Differential Privacy. 3 Credits.

This course provides an introduction to differential privacy, with a focus on algorithmic aspects (rather than statistical or engineering aspects). Specific topics we will cover include: motivation for differential privacy, and different versions of differential privacy (pure, approximate, Renyi, and zero-concentrated in particular); basic mechanisms (Laplace, Gaussian, Discrete Gaussian, and Exponential); composition theorems; basic algorithmic techniques (sparse vector technique, private multiplicative weights, private selection); beyond global sensitivity: local sensitivity, propose-test-release, subsampling; differentially private graph algorithms; lower bounds.

Prerequisite(s): Students may earn credit for EN.601.438 or EN.601.638, but not both.;EN.601.433 OR EN.601.633
Distribution Area: Engineering, Quantitative and Mathematical Sciences

EN.601.440. Web Security. 3 Credits.

This course begins with reviewing basic knowledge of the World Wide Web, and then exploring the central defense concepts behind Web security, such as same-origin policy, cross-origin resource sharing, and browser sandboxing. It will cover the most popular Web vulnerabilities, such as cross-site scripting (XSS) and SQL injection, as well as how to attack and penetrate software with such vulnerabilities. Students will learn how to detect, respond, and recover from security incidents. Newly proposed research techniques will also be discussed.

Prerequisite(s): Students may receive credit for only one of 340/440/640.;EN.601.226 AND EN.601.229 AND EN.601.280

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.441. Blockchains and Cryptocurrencies. 3 Credits.

This course will introduce students to cryptocurrencies and the main underlying technology of Blockchains. The course will start with the relevant background in cryptography and then proceed to cover the recent advances in the design and applications of blockchains. This course should primarily appeal to students who want to conduct research in this area or wish to build new applications on top of blockchains. It should also appeal to those who have a casual interest in this topic or are generally interested in cryptography. Students are expected to have mathematical maturity.

Prerequisite(s): Students may receive credit for only one of EN.600.451 OR EN.601.441 OR EN.601.641;EN.601.226 AND (EN.553.211 OR EN.553.310 OR EN.553.311 OR EN.560.348 OR EN.553.420 OR EN.553.421)

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.442. Modern Cryptography. 3 Credits.

Modern Cryptography includes seemingly paradoxical notions such as communicating privately without a shared secret, proving things without leaking knowledge, and computing on encrypted data. In this challenging but rewarding course we will start from the basics of private and public key cryptography and go all the way up to advanced notions such as zero-knowledge proofs, functional encryption and program obfuscation. The class will focus on rigorous proofs and require mathematical maturity.

Prerequisite(s): Students may receive credit for only one of EN.600.442, EN.601.442, EN.601.642.;(EN.601.230 OR EN.601.231) AND (EN.553.310 OR EN.553.311 OR EN.553.420 OR EN.553.421)

Distribution Area: Engineering, Quantitative and Mathematical Sciences

AS Foundational Abilities: Science and Data (FA2)

EN.601.443. Security & Privacy in Computing. 3 Credits.

Lecture topics will include computer security, network security, basic cryptography, system design methodology, and privacy. There will be a heavy work load, including written homework, programming assignments, exams and a comprehensive final. The class will also include a semester-long project that will be done in teams and will include a presentation by each group to the class.

Prerequisite(s): EN.601.229 (Computer System Fundamentals);Students may receive credit for only one of EN.600.443, EN.601.443, EN.601.643.

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.444. Medical Device Cybersecurity. 3 Credits.

In an increasingly connected healthcare landscape, medical devices have effectively become IT endpoints, often running general-purpose operating systems like Windows or Linux, incorporating cloud microservices, and integrating artificial intelligence to detect, prevent, and improve patient health outcomes. Protecting these devices from cyber threats is not just a technical challenge—it's a matter of patient safety. A security breach in medical devices like pacemakers or infusion pumps can have life-threatening consequences. National and international regulatory bodies, such as the FDA and EU National Competent Authorities (NCAs) and Medical Device Regulation (MDR), know the implications and have provided prescription and guidance emphasizing stringent cybersecurity measures throughout a device's lifecycle, from design and development to postmarket surveillance. The result is a heightened awareness of medical device security and its impact on healthcare delivery, requiring cybersecurity risk management. In particular, focusing on threat modeling, cybersecurity risk assessment, secure design, secure coding practices, vulnerability management and monitoring, software bill of materials, cybersecurity transparency, user labeling, penetration testing, and more. Recommended background: computing systems, operating systems, machine learning & AI.

Prerequisite(s): Students who have taken, or are currently enrolled in EN.601.644 are not eligible to enroll in EN.601.444;EN.601.443 OR EN.601.643

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.445. Practical Cryptographic Systems. 3 Credits.

This semester-long course will teach systems and cryptographic design principles by example: by studying and identifying flaws in widely-deployed cryptographic products and protocols. Our focus will be on the techniques used in practical security systems, the mistakes that lead to failure, and the approaches that might have avoided the problem. We will place a particular emphasis on the techniques of provable security and the feasibility of reverse-engineering undocumented cryptographic systems.

Prerequisite(s): Students may receive credit for only one of EN.601.445 OR EN.601.645, but not both.;EN.601.226 AND EN.601.229

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.446. Sketching and Indexing for Sequences. 3 Credits.

Many of the world's largest and fastest-growing datasets are text, e.g. DNA sequencing data, web pages, logs and social media posts. Such datasets are useful only to the degree we can query, compare and analyze them. Here we discuss two powerful approaches in this area. We will cover sketching, which enables us to summarize very large texts in small structures that allow us to measure the sizes of sets and of their unions and intersections. This in turn allows us to measure similarity and find near neighbors. Second, we will discuss indexing — succinct and compressed indexes in particular — which enables us to efficiently search inside very long strings, especially in highly repetitive texts.

Prerequisite(s): Students can receive credit for EN.601.446 or EN.601.646, but not both;EN.601.226

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.447. Computational Genomics: Sequences. 3 Credits.

Your genome is the blueprint for the molecules in your body. It's also a string of letters (A, C, G and T) about 3 billion letters long. How does this string give rise to you? Your heart, your brain, your health? This, broadly speaking, is what genomics research is about. This course will familiarize you with a breadth of topics from the field of computational genomics. The emphasis is on current research problems, real-world genomics data, and efficient software implementations for analyzing data. Topics will include: string matching, sequence alignment and indexing, assembly, and sequence models. Course will involve significant programming projects.

Prerequisite(s): Students may receive credit for only one of the following: EN.601.447 OR EN.601.647 but not both.;EN.601.220 AND EN.601.226
Distribution Area: Engineering
AS Foundational Abilities: Science and Data (FA2)

EN.601.449. Computational Genomics: Applied Comparative Genomics. 3 Credits.

The goal of this course is to study the leading computational and quantitative approaches for comparing and analyzing genomes starting from raw sequencing data. The course will focus on human genomics and human medical applications, but the techniques will be broadly applicable across the tree of life. The topics will include genome assembly & comparative genomics, variant identification & analysis, gene expression & regulation, personal genome analysis, and cancer genomics. The grading will be based on assignments, a midterm exam, class presentations, and a significant class project. Prerequisites: knowledge of the Unix operating system and programming expertise in a language such as R or Python.

Prerequisite(s): Students may receive credit for only one of EN.600.449, EN.600.649, EN.601.749.
Distribution Area: Engineering
AS Foundational Abilities: Science and Data (FA2)

EN.601.451. Introduction to Computational Immunogenomics. 3 Credits.

Immunology studies defensive mechanisms of living organisms against external threats. Computational immunogenomics is a new field of bioinformatics that develops and applies computational approaches to the study and interpretation of immunological data, seeking to answer questions about adaptive immune responses in humans and important animals. In this course, students will learn how to design, apply, and benchmark algorithms for solving immunogenomics problems. Students may receive credit for only one of EN.601.451, EN.601.651.

Prerequisite(s): Students may receive credit for only one of EN.601.451 OR EN.601.651.;EN.601.220 AND EN.601.226
Distribution Area: Engineering
AS Foundational Abilities: Science and Data (FA2)

EN.601.453. Applications of Augmented Reality. 3 Credits.

This course is designed to expand the student's augmented reality knowledge and introduce relevant topics necessary for developing more meaningful applications and conducting research in this field. The course addresses the fundamental concepts of visual perception and introduces non-visual augmented reality modalities, including auditory, tactile, gustatory, and olfactory applications. The following sessions discuss the importance of integrating user-centered design concepts to design meaningful augmented reality applications. A later module introduces the basic requirements to design and conduct user studies and guidelines on interpreting and evaluating the results from the studies. During the course, students conceptualize, design, implement and evaluate the performance of augmented reality solutions for their use in industrial applications, teaching and training, or healthcare settings. Homework in this course will relate to applying the theoretical methods used for designing, implementing, and evaluating augmented reality applications.

Prerequisite(s): Students may receive credit for only one of EN.601.453 OR EN.601.653, but not both.;EN.601.454 OR EN.601.654
Distribution Area: Engineering
AS Foundational Abilities: Science and Data (FA2)

EN.601.454. Introduction to Augmented Reality. 3 Credits.

This course introduces students to the field of Augmented Reality. It reviews its basic definitions, principles, and applications. The course explains how fundamentals concepts of computer vision are applied for the development of Augmented Reality applications. It then focuses on describing the principal components and particular requirements to implement a solution using this technology. The course also discusses the main issues of calibration, tracking, multi-modal registration, advanced visualization, and display technologies. Homework in this course will relate to the mathematical methods used for calibration, tracking, and visualization in augmented reality.

Prerequisite(s): Students may receive credit for only one of the following: EN.601.454 OR EN.601.654, but not both.;EN.601.220 AND EN.601.226 AND (AS.110.201 OR AS.110.212 OR EN.553.291 OR EN.553.295)
Distribution Area: Engineering
AS Foundational Abilities: Science and Data (FA2)

EN.601.455. Computer Integrated Surgery I. 4 Credits.

This course focuses on computer-based techniques, systems, and applications exploiting quantitative information from medical images and sensors to assist clinicians in all phases of treatment from diagnosis to preoperative planning, execution, and follow-up. It emphasizes the relationship between problem definition, computer-based technology, and clinical application and includes a number of guest lectures given by surgeons and other experts on requirements and opportunities in particular clinical areas. Recommended Course Background: EN.601.220, EN.601.457, EN.601.461, image processing.

Prerequisite(s): Students may receive credit for only one of EN.600.445, EN.600.645, EN.601.455, EN.601.655.;EN.600.226/EN.601.226 AND (AS.110.201 OR AS.110.212 OR EN.553.291 OR EN.553.295) or permission of the instructor.
Distribution Area: Engineering
AS Foundational Abilities: Science and Data (FA2)

EN.601.456. Computer Integrated Surgery II. 3 Credits.

This weekly lecture/seminar course addresses similar material to EN.601.455, but covers selected topics in greater depth. In addition to material covered in lectures/seminars by the instructor and other faculty, students are expected to read and provide critical analysis/presentations of selected papers in recitation sessions. Students taking this course are required to undertake and report on a significant term project under the supervision of the instructor and clinical end users. Typically, this project is an extension of the term project from EN.601.455, although it does not have to be. Grades are based both on the project and on classroom recitations. Students who wish to use this course to satisfy the "Team" requirement should register for EN.601.496 instead. Students wishing to attend the weekly lectures as a 1-credit seminar should sign up for EN.601.356.

Prerequisite(s): Students may receive credit for only one of EN.600.446, EN.600.646, EN.601.456, EN.601.496, OR EN.601.656.;EN.601.455 or EN.601.655 or permission

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.457. Computer Graphics. 3 Credits.

This course introduces computer graphics techniques and applications, including image processing, rendering, modeling and animation.

Prerequisite(s): Students may receive credit for only one of the following: EN.601.457 OR EN.601.657, but not both.;EN.601.220 AND EN.601.226 AND (AS.110.201 OR AS.110.212 OR EN.553.291 OR EN.553.295)

Distribution Area: Engineering, Quantitative and Mathematical Sciences

AS Foundational Abilities: Science and Data (FA2)

EN.601.461. Computer Vision. 3 Credits.

This course provides an overview of fundamental methods in computer vision from a computational perspective. Methods studied include: camera systems and their modelling, computation of 3-D geometry from binocular stereo, motion, and photometric stereo, and object recognition, image segmentation, and activity analysis. Elements of machine vision and biological vision are also included.

Prerequisite(s): Students may receive credit for only one of EN.600.361, EN.600.461, EN.600.661, EN.601.461, EN.601.661.;(EN.553.310 OR EN.553.311 OR ((EN.553.420 OR EN.553.421) AND (EN.553.430 OR EN.553.431)) AND (AS.110.201 OR AS.110.212 OR EN.553.291 OR EN.553.295)) AND (EN.500.112 OR EN.500.113 OR EN.500.114 OR EN.601.220 AS.250.205)

Distribution Area: Engineering, Quantitative and Mathematical Sciences

AS Foundational Abilities: Science and Data (FA2)

EN.601.463. Algorithms for Sensor-Based Robotics. 3 Credits.

This course surveys the development of robotic systems for navigating in an environment from an algorithmic perspective. It will cover basic kinematics, configuration space concepts, motion planning, and localization and mapping. It will describe these concepts in the context of the ROS software system, and will present examples relevant to mobile platforms, manipulation, robotics surgery, and human-machine systems.

Prerequisite(s): Students may receive credit for only one of EN.600.336, EN.600.436, EN.600.636, EN.601.463, EN.601.663.;(AS.110.201 OR AS.110.212 OR EN.553.291 OR EN.553.295) AND (AS.110.202 OR AS.110.211) AND EN.601.226 AND (EN.553.211 OR EN.553.310 OR EN.553.311 OR EN.553.420 OR EN.553.421)

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.464. Artificial Intelligence. 3 Credits.

The course situates the study of Artificial Intelligence (AI) first in the broader context of Cognitive Science (i.e., human intelligence) and then treats in-depth principles and methods for reasoning, planning, and learning, including both conventional methods and recent deep learning approaches. The class is recommended for all scientists and engineers with a genuine curiosity about how to build an AI system (in particular, an intelligent agent) that can learn, reason about, and interact with the world and other agents. Strong programming skills and a solid mathematical foundation are expected. Students will be asked to complete both programming assignments and writing assignments. Students can only receive credit for one of 601.464/664.

Prerequisite(s): EN.600.226/EN.601.226;Students may receive credit for only one of EN.600.335, EN.600.435, EN.601.464, EN.601.664.

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.465. Natural Language Processing. 4 Credits.

An in-depth introduction to core techniques for analyzing, transforming, and generating human language. The course spans linguistics, modeling, algorithms, and applications. (1) How should linguistic structure and meaning be represented (e.g., trees, morphemes, ?-terms, vectors)? (2) How can we formally model the legal structures and their probabilities (e.g., grammars, automata, features, log-linear models, recurrent neural nets, Transformers)? (3) What algorithms can estimate the parameters of these models (e.g., gradient descent, EM) and efficiently identify probable structures (e.g., dynamic programming, beam search)? (4) Finally, what kinds of systems can be built with these techniques and how are they constructed and evaluated in practice? Detailed assignments guide students through many details of implementing core NLP methods.

The course proceeds from first principles, although prior exposure to AI, statistics, ML, or linguistics can be helpful. Prerequisite: Data Structures and basic familiarity with Python, partial derivatives, matrix multiplication and probabilities.

Prerequisite(s): EN.600.226/EN.601.226;Students may receive credit for only one of EN.600.465, EN.601.465, EN.601.665.

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.466. Information Retrieval and Web Agents. 3 Credits.

An in-depth, hands-on study of current information retrieval techniques and their application to developing intelligent WWW agents. Topics include a comprehensive study of current document retrieval models, mail/news routing and filtering, document clustering, automatic indexing, query expansion, relevance feedback, user modeling, information visualization and usage pattern analysis. In addition, the course explores the range of additional language processing steps useful for template filling and information extraction from retrieved documents, focusing on recent, primarily statistical methods. The course concludes with a study of current issues in information retrieval and data mining on the World Wide Web. Topics include web robots, spiders, agents and search engines, exploring both their practical implementation and the economic and legal issues surrounding their use. Recommended Course Background: EN.601.226

Prerequisite(s): Students can receive credit for EN.601.466 or EN.601.666, but not both;EN.600.226 OR EN.601.226

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.467. Introduction to Human Language Technology. 3 Credits.

This course gives an overview of basic foundations and applications of human language technology, such as: morphological, syntactic, semantic, and pragmatic processing; machine learning; signal processing; speech recognition; speech synthesis; information retrieval; text classification; topic modelling; information extraction; knowledge representation; machine translation; dialog systems; etc. Pre-req: EN.601.226 Data Structures; knowledge of Python recommended.

Prerequisite(s): Students can receive credit for only one of 601.467/601.647;EN.601.226 OR EN.600.226

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.468. Machine Translation. 3 Credits.

Google translate can instantly translate between any pair of over fifty human languages (for instance, from French to English). How does it do that? Why does it make the errors that it does? And how can you build something better? Modern translation systems learn to translate by reading millions of words of already translated text, and this course will show you how they work. The course covers a diverse set of fundamental building blocks from linguistics, machine learning, algorithms, data structures, and formal language theory, along with their application to a real and difficult problem in artificial intelligence.

Prerequisite(s): Students may receive credit for only one of EN.600.468, EN.601.468, EN.601.668.;EN.601.226 AND EN.553.211 OR EN.553.310 OR EN.553.311 OR ((EN.553.420 OR EN.553.421) AND (EN.553.430 OR EN.553.431)))

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.469. AI Safety, Alignment, & Governance. 3 Credits.

This course will focus on the alignment and governance challenges posed by advanced frontier/general purpose AI models: why these models may behave in ways that pose significant risk to human welfare and what technical and governance approaches might mitigate these risks. We'll begin the course studying general results from alignment and governance in human normative systems such as markets, politics, norms and laws. We'll pay special attention to risks arising from agentic AI. We'll then look at current technical and position papers in various topics in AI safety and alignment. Topics could include: RLHF, constitutional AI, red-teaming, safety evaluation methods, red lines, jail-breaking, prompt injection, over-optimization, and open-source debates. We'll conclude with discussion of regulatory frameworks such as regulatory markets, registration of frontier models, international governance organizations, registration of AI agents and legal personhood for AI agents. This is a paper-reading class.

Prerequisite(s): Students may only receive credit for EN.601.469 OR EN.601.669.;En.601.474 OR En.601.674 OR En.601.475 OR En.601.675 OR En.601.482 OR En.601.682 OR En.601.486 OR En.601.686

Distribution Area: Engineering

EN.601.470. Artificial Agents. 3 Credits.

This course covers a number of topics explored in introductory AI, such as knowledge representation, reasoning, and natural language understanding. Unlike introductory AI, we will pursue these topics based on the transformer neural architecture. We will motivate the material through interacting with agents in games: how to build models that understand user commands, how to generate responses back to a user, and how to reason about a synthetic environment to determine a course of action. Assignments will include programming, presentations on readings, and written summaries of readings.

Prerequisite(s): (EN.601.475 OR EN.601.675) OR (EN.601.482 OR EN.601.682) OR (EN.601.488 OR EN.601.688) OR (EN.601.486 OR EN.601.686)

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.471. Natural Language Processing: Self-Supervised Models. 3 Credits.

The rise of massive self-supervised (pre-trained) models have transformed various data-driven fields such as natural language processing (NLP). In this course, students will gain a thorough introduction to self-supervised learning techniques for NLP applications. Through lectures, assignments, and a final project, students will learn the necessary skills to design, implement, and understand their own self-supervised neural network models, using the Pytorch framework. Students should have familiarity with Python/PyTorch.

Prerequisite(s): Students may receive credit for EN.601.471 or EN.601.671, but not both.;EN.601.226 AND (EN.553.211 OR EN.553.310 OR EN.553.311 OR EN.553.420 OR EN.553.421) AND (AS.110.201 OR AS.110.212 OR EN.553.291 OR EN.553.295)

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.472. Natural Language Processing for Computational Social Science. 3 Credits.

[Alt. title: Analyzing Text as Data] Vastly available digitized text data has created new opportunities for understanding social phenomena. Relatedly, social issues like toxicity, discrimination, and propaganda frequently manifest in text, making text analyses critical for understanding and mitigating them. In this course, we will centrally explore: how can we use NLP as a tool for understanding society? Students will learn core and recent advances in text-analysis methodology, building from word-level metrics to embeddings and language models as well as incorporating statistical methods such as time series analyses and causal inference. Students may receive credit for EN.601.472 or EN.601.672, but not both.

Prerequisite(s): Students may receive credit for only one of the following: EN.601.472 OR EN.601.672;((EN.601.465 OR EN.601.665) OR (EN.601.467 OR EN.601.667) OR (EN.601.468 OR EN.601.668))

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.473. Cognitive Artificial Intelligence. 3 Credits.

Humans, even young children, can learn, model, and reason about the world and other people in a fast, robust, and data efficient way. This course will discuss the principles of human cognition, how we can use machine learning and AI models to computationally capture these principles, and how these principles can help us build better AI. Topics will include (but are not limited to) Bayesian concept learning, probabilistic programming, intuitive physics, decision-making, Theory of Mind, pragmatics, and value alignment. Strongly recommended: a prior course in machine learning or artificial intelligence.

Prerequisite(s): Students who have taken EN.601.673 are not eligible to enroll in EN.601.473.;(((EN.553.420 OR EN.553.421) AND (EN.553.430 OR EN.553.431)) OR (EN.553.211 OR EN.553.310 OR EN.553.311) AND (AS.110.201 OR AS.110.212 OR EN.553.291 OR EN.553.295) AND (EN.500.112 OR EN.500.113 OR EN.500.114 OR EN.601.220 OR AS.250.205)))

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.474. ML: Learning Theory. 3 Credits.

This machine learning course will provide a formal and in-depth coverage of topics in statistical and computational learning theory. We will revisit popular machine learning algorithms and understand their performance in terms of the size of the data (sample complexity), memory needed (space complexity), as well as the overall runtime (computational or iteration complexity). We will cover topics including PAC learning, uniform convergence, VC dimension, Rademacher complexity, algorithmic stability, kernel methods, online learning and reinforcement learning, as well as introduce students to current topics in large-scale machine-learning and randomized projections. General focus will be on combining methodology with theoretical and computational foundations.

Prerequisite(s): (AS.110.202 OR AS.110.211) AND (((EN.553.420 OR EN.553.421) AND (EN.553.430 OR EN.553.431)) OR (EN.553.211 OR EN.553.310 OR EN.553.311) AND (AS.110.201 OR AS.110.212 OR EN.553.291 OR EN.553.295)) AND (EN.500.112 OR EN.500.113 OR EN.500.114) OR (EN.601.220 OR AS.250.205 OR EN.580.200 OR EN.601.107)

Distribution Area: Engineering, Quantitative and Mathematical Sciences

AS Foundational Abilities: Science and Data (FA2)

EN.601.475. Machine Learning. 3 Credits.

Machine learning is subfield of computer science and artificial intelligence, whose goal is to develop computational systems, methods, and algorithms that can learn from data to improve their performance. This course introduces the foundational concepts of modern Machine Learning, including core principles, popular algorithms and modeling platforms. This will include both supervised learning, which includes popular algorithms like SVMs, logistic regression, boosting and deep learning, as well as unsupervised learning frameworks, which include Expectation Maximization and graphical models. Homework assignments include a heavy programming components, requiring students to implement several machine learning algorithms in a common learning framework. Additionally, analytical homework questions will explore various machine learning concepts, building on the pre-requisites that include probability, linear algebra, multi-variate calculus and basic optimization. Students in the course will develop a learning system for a final project.

Prerequisite(s): Students may receive credit for only one of EN.600.475, EN.601.475, EN.601.675.;Linear Algebra, Probability, Statistics, Calc III, and Intro Computing/Programming - (AS.110.202 OR AS.110.211) AND (EN.553.211 OR EN.553.310 OR EN.553.311 OR ((EN.553.420 or EN.553.421) AND (EN.553.430 OR EN.553.431)) AND (AS.110.201 OR AS.110.212 OR EN.553.291 OR EN.553.295) AND (EN.500.112 OR EN.500.113 OR EN.500.114 OR (EN.601.220 OR EN.600.120) OR AS.250.205 OR EN.580.200 OR (EN.600.107 OR EN.601.107)).

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.477. Causal Inference. 3 Credits.

"Big data" is not necessarily "high quality data." Systematically missing records, unobserved confounders, and selection effects present in many datasets make it harder than ever to answer scientifically meaningful questions. This course will teach mathematical tools to help you reason about causes, effects, and bias sources in data with confidence. We will use graphical causal models, and potential outcomes to formalize what causal effects mean, describe how to express these effects as functions of observed data, and use regression model techniques to estimate them. We will consider techniques for handling missing values, structure learning algorithms for inferring causal directionality from data, and connections between causal inference and reinforcement learning. Pre-requisites: familiarity with the R programming language, multivariate calculus, basics of linear algebra and probability.

Prerequisite(s): Students may receive credit for only one of EN.600.477, EN.600.677, EN.601.477, EN.601.677.;EN.601.475 OR (EN.553.211 OR EN.553.311 OR EN.553.420 OR EN.553.421) AND (AS.110.202 OR AS.110.211) or permission of instructor.

Distribution Area: Engineering, Quantitative and Mathematical Sciences

AS Foundational Abilities: Science and Data (FA2)

EN.601.479. Machine Learning: Reinforcement Learning. 3 Credits.

Tremendous success of reinforcement learning (RL) in a variety of settings from AlphaGo to LLMs makes it a critical area to study. This course will study classical aspects of RL as well as its modern counterparts. Topics will include Markov Decision Processes, dynamic programming, model-based and model-free RL, temporal difference learning, Monte Carlo methods, multi-armed bandits, policy optimization and other methods.

Prerequisite(s): Students who have taken, or are currently enrolled in EN.601.679, are not eligible to take EN.601.479.;(((EN.601.464 OR EN.601.664) OR (EN.601.475 OR EN.601.675) OR (EN.601.482 OR EN.601.682)) AND (AS.110.201 OR AS.110.212 OR EN.553.291 OR EN.553.295) AND (EN.553.211 OR EN.553.310 OR EN.553.311 OR EN.553.420 OR EN.553.421))

Distribution Area: Engineering

EN.601.482. Machine Learning: Deep Learning. 4 Credits.

Deep learning (DL) has emerged as a powerful tool for solving data-intensive learning problems such as supervised learning for classification or regression, dimensionality reduction, and control. As such, it has a broad range of applications including speech and text understanding, computer vision, medical imaging, and perception-based robotics. The goal of this course is to introduce the basic concepts of deep learning (DL). The course will include a brief introduction to the basic theoretical and methodological underpinnings of machine learning, commonly used architectures for DL, DL optimization methods, DL programming systems, and specialized applications to computer vision, speech understanding, and robotics. Students will be expected to solve several DL problems on standardized data sets, and will be given the opportunity to pursue team projects on topics of their choice. Strongly recommended courses: Python, Machine Learning, Statistics, Calc III

Prerequisite(s): Students can receive credit for EN.601.482 or EN.601.682, but not both;EN.601.226 AND (AS.110.201 OR AS.110.212 OR EN.553.291 OR EN.553.295) AND (EN.553.211 OR EN.553.310 OR EN.553.311 OR EN.553.420 OR EN.553.421) AND (AS.110.107 OR AS.110.109 OR AS.110.113); Strongly recommended courses: Python, Machine Learning, Statistics, Calc III

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.484. Explainable AI Design & Human-AI Interaction. 3 Credits.

This is a design course. Increasing the trustworthiness of machine learning solutions has emerged as an important research area. One approach to trustworthy machine learning is explainable and/or interpretable machine learning, which attempts to reveal the working mechanisms of a machine learning system. However, other than on-task performance, explainability is not a property of machine learning algorithms, but an affordance: a relationship between explanation model and the target users in their context. Successful development of machine learning solutions that afford explainability thus requires understanding of techniques beyond pure machine learning. In this course, we will first review the basics of machine learning and human-centered design. Then, we will introduce several techniques to explain machine learning models and/or make them interpretable, and through hands-on sessions and case studies, will investigate how these techniques affect human-AI interaction. In addition to individual homework assignments, students will work in groups to design, justify, implement, and test an explainable machine learning algorithm for a problem of their choosing. Additional Recommended (601.454/654, 601.290, 601.490/690 or 601.491/691) and 601.477/677.

Prerequisite(s): Students may receive credit for EN.601.484 or EN.601.684, but not both.;(EN.601.476 OR EN.601.676) OR (EN.601.464 OR EN.601.664) OR (EN.601.482 OR EN.601.682)

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.485. Probabilistic Models of the Visual Cortex. 3 Credits.

The course gives an introduction to computational models of the mammalian visual cortex. It covers topics in low-, mid-, and high-level vision. It briefly discusses the relevant evidence from anatomy, electrophysiology, imaging (e.g., fMRI), and psychophysics. It concentrates on mathematical modelling of these phenomena taking into account recent progress in probabilistic models of computer vision and developments in machine learning, such as deep networks. Programming experience (Python preferred).

Prerequisite(s): Students who have taken AS.050.375/AS.050.675/ EN.601.685 are not eligible to take EN.601.485.;AS.110.106 OR AS.110.108

Distribution Area: Quantitative and Mathematical Sciences

AS Foundational Abilities: Science and Data (FA2)

EN.601.486. Machine Learning: Artificial Intelligence System Design & Development. 3 Credits.

Advances in Artificial intelligence have opened new opportunities for developing systems to aid in numerous areas of society. In order for AI systems to succeed in making constructive and positive changes, we must consider their impact on everyday life. Specifically, AI system designers must evaluate the overall capabilities of the system, consider the resulting human-AI interactions, and ensure that the system behaves in a responsible and ethical manner. In this project-based course you will work in teams of 3-5 students to 1) Identify a need with high-impact implications on everyday life; 2) Articulate principles of Responsible AI relevant to the intended application, 3) Conceptualize and design an AI system targeting this need, and 4) Develop the AI system by refining a demo-able prototype based on feedback received during course presentations. Additionally, we will discuss potential ethical issues that can arise in AI and how to develop Responsible AI principles. Coursework will consist of writing assignments, project presentations, and a project demonstration. Recommended background: Python programming, EN.601.290 or EN.601.454/654 or EN.601.490/690 or EN.601.491/691.

Prerequisite(s): (EN.601.475 OR EN.601.675) OR (EN.601.464 OR EN.601.664) OR (EN.601.482 OR EN.601.682)

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.487. Machine Learning: Coping with Non-Stationary Environments and Errors. 3 Credits.

This course teaches machine learning methods that 1) consider data distribution shift and 2) represent and quantify the model uncertainty in a principled way. The topics we will cover include machine learning techniques that deal with data distribution shift, including domain adaptation, domain generalization, and distributionally robust learning techniques, and various uncertainty quantification methods, including Bayesian methods, conformal prediction methods, and model calibration methods. We will introduce these topics in the context of building trustworthy machine learning solutions to safety-critical applications and socially-responsible applications. For example, a typical application is responsible decision-making under uncertainty in non-stationary environments. So we will also introduce concepts like fair machine learning and learning under safety constraints, and discuss how robust and uncertainty-aware learning techniques contribute to such more desired systems. Students will learn the state-of-the-art methods in lectures, test their understanding in homeworks, and apply these methods in a project.

Prerequisite(s): Students who have taken, or are currently enrolled in EN.601.687 OR EN.601.787, are not eligible to take EN.601.487.;EN.601.475 OR EN.601.675

Distribution Area: Engineering

EN.601.489. Human-in-the-Loop Machine Learning. 4 Credits.

Machine learning (ML) is being deployed in increasingly consequential tasks, such as healthcare and autonomous driving. For the foreseeable future, successfully deploying ML in such settings will require close collaboration and integration with humans, whether they be users, designers, engineers, policy-makers, etc. This course will look at how humans can be incorporated into the foundations of ML in a principled way. The course will be broken down into three parts: demonstration, collaboration, and oversight. Demonstration is about how machines can learn from 'observing' humans—such as learning to drive a car from data collected while humans drive. In this setting, the human is assumed to be strictly better than the machine and so the primary goal is to transmit the human's knowledge and abilities into the ML model. The second part, collaboration, is about when humans and models are near equals in performance but not in abilities. A relevant setting is AI-assisted healthcare: perhaps a human radiologist and ML model are good at diagnosing different kinds of diseases. Thus we will look at methodologies that allow machines to 'ask for help' when they are either unconfident in their own performance and/or think the human can better carry out the task. The course will close with the setting in which machines are strictly better at a task than humans are, but we still wish to monitor them to ensure safety and alignment with our goals (oversight). Assessment will be done with homework, quizzes, and a final project.

Prerequisite(s): Students who have taken or are enrolled in EN.601.689 are not eligible to take EN.601.489.;EN.601.475 OR EN.601.675

Distribution Area: Engineering

EN.601.490. Introduction to Human-Computer Interaction. 3 Credits.

This course is designed to introduce undergraduate and graduate students to design techniques and practices in human-computer interaction (HCI), the study of interactions between humans and computing systems. Students will learn design techniques and evaluation methods, as well as current practices and exploratory approaches, in HCI through lectures, readings, and assignments. Students will practice various design techniques and evaluation methods through hands-on projects focusing on different computing technologies and application domains. This course is intended for undergraduate and graduate students in Computer Science/Cognitive Science/Psychology. Interested students from different disciplines should contact the instructor before enrolling in this course. Recommended Background: Basic programming skills.

Prerequisite(s): Students can receive credit for either EN.601.490 or EN.601.690, but not both.

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.491. Human-Robot Interaction. 3 Credits.

This course is designed to introduce advanced students to research methods and topics in human-robot interaction (HRI), an emerging research area focusing on the design and evaluation of interactions between humans and robotic technologies. Students will (1) learn design principles for building and research methods of evaluating interactive robot systems through lectures, readings, and assignments, (2) read and discuss relevant literature to gain sufficient knowledge of various research topics in HRI, and (3) work on a substantial project that integrates the principles, methods, and knowledge learned in this course.

Prerequisite(s): Students can receive credit for EN.601.491 or EN.601.691, but not both;EN.601.220/EN.600.120 AND EN.601.226/ EN.600.226

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.493. Accessible Computing. 3 Credits.

This course is designed to introduce students to the principles, challenges, and opportunities in designing computing systems that are accessible to people with diverse abilities. Students will learn about assistive technologies, inclusive design methodologies, and the incorporation of accessibility in computer applications through lectures, readings, and projects.

Prerequisite(s): Students may only receive credit for EN.601.493 OR EN.601.693.;EN.601.290 OR EN.601.490 OR EN.601.690

Distribution Area: Engineering

EN.601.496. Computer Integrated Surgery II - Teams. 3 Credits.

This weekly lecture/seminar course addresses similar material to 600.455, but covers selected topics in greater depth. In addition to material covered in lectures/seminars by the instructor and other faculty, students are expected to read and provide critical analysis/presentations of selected papers in recitation sessions. Students taking this course are required to undertake and report on a significant term project in teams of at least 3 students, under the supervision of the instructor and clinical end users. Typically, this project is an extension of the term project from 600.455, although it does not have to be. Grades are based both on the project and on classroom recitations. Students who prefer to do individual projects must register for EN.601.456 instead.

Prerequisite(s): Students may receive credit for only one of EN.601.456, EN.601.496, OR EN.601.656;EN.601.455 or permission

Distribution Area: Engineering

AS Foundational Abilities: Science and Data (FA2)

EN.601.501. Computer Science Workshop. 1 - 3 Credits.

An applications-oriented, computer science project done under the supervision and with the sponsorship of a faculty member in the Department of Computer Science. Computer Science Workshop provides a student with an opportunity to apply theory and concepts of computer science to a significant project of mutual interest to the student and a Computer Science faculty member. Permission to enroll in CSW is granted by the faculty sponsor after his/her approval of a project proposal from the student. Interested students are advised to consult with Computer Science faculty members before preparing a Computer Science Workshop project proposal.

Prerequisite(s): You must request Customized Academic Learning using the Customized Academic Learning form found in Student Self-Service: Registration > Online Forms.

EN.601.503. Independent Study. 1 - 3 Credits.

Individual guided study for undergraduate students under the direction of a faculty member in the department. The program of study, including the credit to be assigned, must be worked out in advance between the student and the faculty member involved. Permission required.

Prerequisite(s): You must request Customized Academic Learning using the Customized Academic Learning form found in Student Self-Service: Registration > Online Forms.

EN.601.507. Undergraduate Research. 1 - 3 Credits.

Individual research for undergraduates under the direction of a faculty member in the department. The program of research, including the credit to be assigned, must be worked out in advance between the student and the faculty member involved. Permission required.

Prerequisite(s): You must request Customized Academic Learning using the Customized Academic Learning form found in Student Self-Service: Registration > Online Forms.

EN.601.509. Computer Science Internship. 1 Credit.

Individual work in the field with a learning component, supervised by a faculty member in the department. The program of study and credit assigned must be worked out in advance between the student and the faculty member involved. As a rule of thumb, 40 hours of work is equivalent to one credit. Permission required.

Prerequisite(s): You must request Customized Academic Learning using the Customized Academic Learning form found in Student Self-Service: Registration > Online Forms.

EN.601.513. Group Undergraduate Project. 1 - 3 Credits.

Independent learning and application for undergraduates under the direction of a faculty member in the department. This course has a regular project group meeting that students are expected to attend. The individual project contributions, including the credit to be assigned, must be worked out in advance between the student and the faculty member involved.

Prerequisite(s): You must request Customized Academic Learning using the Customized Academic Learning form found in Student Self-Service: Registration > Online Forms.

EN.601.517. Group Undergraduate Research. 1 - 3 Credits.

Independent research for undergraduates under the direction of a faculty member in the department. This course has a weekly research group meeting that students are expected to attend. The program of research, including the credit to be assigned, must be worked out in advance between the student and the faculty member involved.

Prerequisite(s): You must request Customized Academic Learning using the Customized Academic Learning form found in Student Self-Service: Registration > Online Forms.

EN.601.519. Senior Honors Thesis. 3 Credits.

The student will undertake a substantial independent research project under the supervision of a faculty member, potentially leading to the notation "Departmental Honors with Thesis" on the final transcript. Students are expected to enroll in both semesters of this course during their senior year. Project proposals must be submitted and accepted in the preceding spring semester (junior year) before registration. Students will present their work publicly before April 1st of senior year. They will also submit a first draft of their project report (thesis document) at that time. Faculty will meet to decide if the thesis will be accepted for honors. Computer science majors only. Students should have a 3.5 GPA in computer science courses at the end of their junior year and permission of faculty sponsor.

Prerequisite(s): You must request Customized Academic Learning using the Customized Academic Learning form found in Student Self-Service: Registration > Online Forms.

EN.601.520. Senior Honors Thesis. 1 - 3 Credits.

For computer science majors only, a continuation of EN.601.519.

Recommended Course Background: EN.601.519

Prerequisite(s): You must request Customized Academic Learning using the Customized Academic Learning form found in Student Self-Service: Registration > Online Forms.

EN.601.556. Senior Thesis In CIS. 3 Credits.

The student will undertake a substantial independent research project in the area of computer-integrated surgery, under joint supervision of a WSE faculty adviser and a clinician or clinical researcher at the Johns Hopkins Medical School.

Prerequisite(s): You must request Customized Academic Learning using the Customized Academic Learning form found in Student Self-Service: Registration > Online Forms.;EN.600.445 or permission of instructor.

EN.601.604. Brain & Computation. 1 Credit.

Computational and network aspects of the brain are explored. The topics covered include structure, operation and connectivity of neurons, general network structure of the neural system, and the connectivity constraints imposed by pre- and post-natal neural development and the desirability of network consistency within a species. Both discrete and continuous aspects of neural computation are covered. Precise mathematical tools and analyses such as logic design, transient and steady state behavior of linear systems, and time and connectivity randomization are discussed.

The concepts are illustrated with several applications. Memory formation from the synaptic level to the high level constructs are explored. Students are not expected to master any of the mathematical techniques but are expected to develop a strong qualitative appreciation of their power. Cerebellum, which has a simple network connectivity, will be covered as a typical system. Recommended course background: linear algebra, differential equations, probability, and algorithms.

Prerequisite(s): Students can receive credit for EN.601.404 or EN.601.604, but not both

Distribution Area: Engineering

EN.601.613. Software Defined Networks. 3 Credits.

Software-Defined Networks (SDN) enable programmability of data networks and hence rapid introduction of new services. They use software-based controllers to communicate with underlying hardware infrastructure and direct traffic on a network. This model differs from that of traditional networks, which use dedicated hardware devices (i.e., routers and switches) to control network traffic. This technology is becoming a key part of web scale networks (at companies like Google and Amazon) and 5G/6G networks. Its importance will keep on growing. Many of today's services and applications, especially when they involve the cloud, could not function without SDN. SDN allows data to move easily between distributed locations, which is critical for cloud applications. A major focus will be on how this technology will be used in 5G and 6G Networks. The course will cover basics of SDN, ongoing research in this area, and the industrial deployments. Required Course Background: computer networks.

Prerequisite(s): Students can receive credit for EN.601.413 or EN.601.613, but not both.

Distribution Area: Engineering

EN.601.614. Computer Networks. 3 Credits.

Topics covered will include applications layer protocols (e.g. HTTP, FTP, SMTP), transport layer protocols (UDP, TCP), network layer protocols (e.g. IP, ICMP), link layer protocols (e.g. Ethernet) and wireless protocols (e.g. IEEE 802.11). The course will also cover routing protocols such as link state and distance vector, multicast routing, and path vector protocols (e.g. BGP). The class will examine security issues such as firewalls and denial of service attacks. We will also study DNS, Web caching and CDNS, peer to peer, and protocol tunneling. Finally, we will explore security protocols (e.g. TLS, SSH, IPsec), as well as some basic cryptography necessary to understand these. Grading will be based on hands-on programming assignments, homework and two exams. Required course background: C/C++ programming and data structures, or permission.

Prerequisite(s): Students can only receive credit for EN.601.414 or EN.601.614, but not both.

Distribution Area: Engineering

EN.601.615. Databases. 3 Credits.

Same material as 601.415, for graduate students. Introduction to database management systems and database design, focusing on the relational and object-oriented data models, query languages and query optimization, transaction processing, parallel and distributed databases, recovery and security issues, commercial systems and case studies, heterogeneous and multimedia databases, and data mining. (www.cs.jhu.edu/~yarowsky/cs415.html) Required course background: Data Structures

Prerequisite(s): Students may receive credit for only one of EN.600.315, EN.600.415, EN.601.315, EN.601.415, EN.601.615.

Distribution Area: Engineering

EN.601.618. Operating Systems. 3 Credits.

Same material as 601.418, for graduate students. This course covers fundamental topics related to operating systems theory and practice. Topics include processor management, storage management, concurrency control, multi-programming and processing, device drivers, operating system components (e.g., file system, kernel), modeling and performance measurement, protection and security, and recent innovations in operating system structure. Course work includes the implementation of operating systems techniques and routines, and critical parts of a small but functional operating system. Required course background: Data Structures & Computer System Fundamentals

Prerequisite(s): Students may receive credit for only one of EN.600.318, EN.600.418, EN.601.318, EN.601.418, EN.601.618.

Distribution Area: Engineering

EN.601.619. Cloud Computing. 3 Credits.

Clouds host a wide range of the applications that we rely on today. In this course, we study common cloud applications, traffic patterns that they generate, critical networking infrastructures that support them, and core networking and distributed systems concepts, algorithms, and technologies used inside clouds. We will also study how today's application demand is influencing the network's design, explore current practice, and how we can build future's networked infrastructure to better enable both efficient transfer of big data and low-latency requirements of real-time applications. The format of this course will be a mix of lectures, discussions, assignments, and a project designed to help students practice and apply the theories and techniques covered in the course.

Prerequisites: EN.601.226 or permission. Students can only receive credit for one of 601.419/619. Recommended: a course in operating systems, networks or systems programming.

Prerequisite(s): Students may earn credit for EN.601.419 or EN.601.619, but not both.

Distribution Area: Engineering

EN.601.620. Parallel Computing for Data Science. 3 Credits.

This course studies parallelism in data science, drawing examples from data analytics, statistical programming, and machine learning. It focuses mostly on the Python programming ecosystem but will use C/C++ to accelerate Python and Java to explore shared-memory threading. It explores parallelism at all levels, including instruction level parallelism (pipelining and vectorization), shared-memory multicore, and distributed computing. Concepts from computer architecture and operating systems will be developed in support of parallelism, including Moore's law, the memory hierarchy, caching, processes/threads, and concurrency control. The course will cover modern data-parallel programming frameworks, including Dask, Spark, Hadoop!, and Ray. The course will not cover GPU deep-learning frameworks nor CUDA. The course is suitable for second-year undergraduate CS majors and graduate students from other science and engineering disciplines that have prior programming experience. Required course background: Data Structures, Computer System Fundamentals, and familiarity with Python.

Prerequisite(s): Students may receive credit for only one of EN.601.320, EN.601.420, OR EN.601.620.

Distribution Area: Engineering

EN.601.621. Object Oriented Software Engineering. 3 Credits.

Same material as EN.601.421, for graduate students. This course covers object-oriented software construction methodologies and their application. The main component of the course is a large team project on a topic of your choosing. Course topics covered include object-oriented analysis and design, UML, design patterns, refactoring, program testing, code repositories, team programming, and code reviews. Required course background: intermediate programming, data structures, and experience in mobile or web app development. Students may receive credit for only one of 601.421/621.

Prerequisite(s): Students may receive credit for only one of EN.600.321, EN.600.421, EN.601.421, EN.601.621.

Distribution Area: Engineering

EN.601.622. Software Testing & Debugging. 3 Credits.

Studies show that testing can account for over 50% of software development costs. This course presents a comprehensive study of software testing, principles, methodologies, tools, and techniques. Topics include testing principles, coverage (graph coverage, logic coverage, input space partitioning, and syntax-based coverage), unit testing, higher-order testing (integration, system-level, acceptance), testing approaches (white-box, black-box, grey-box), regression testing, debugging, delta debugging, and several specific types of functional and non-functional testing as schedule/interest permits (GUI testing, usability testing, security testing, load/performance testing, A/B testing etc.). For practical topics, state-of-the-art tools/techniques will be studied and utilized. Required course background: significant mobile or web app development.

Prerequisite(s): Students can only take EN.601.422 or EN.601.622, but not both.

Distribution Area: Engineering

EN.601.624. Reliable Software Systems. 3 Credits.

Reliability is an essential quality requirement for all artifacts operating in the real-world, ranging from bridges, cars to power grids. Software systems are no exception. In this computing age when software is transforming even traditional mission-critical artifacts, making sure the software we write is reliable becomes ever more important. This course exposes students to the principles and techniques in building reliable systems. We will study a set of systematic approaches to make software more robust. These include but are not limited to static analysis, testing framework, model checking, symbolic execution, fuzzing, and formal verification. In addition, we will cover the latest research in system reliability. Recommended course background: EN.601.220 AND EN.601.628.

Prerequisite(s): Students may receive credit for EN.601.424 OR EN.601.624, but not both.

EN.601.625. Software System Design. 3 Credits.

This course introduces modern software systems design, with an emphasis on how to design large-scale systems, assess common system design trade-offs, and tackle system design challenges. It covers non-functional requirements, API design, distributed systems concepts, modern software building blocks (e.g., load balancers, caches, containers, etc.). Additionally, it includes case studies of common system design problems, some drawn from interview questions. Ultimately, this course helps learners become better software engineers. Required course background: EN.601.315/415/615 or EN.601.280 or EN.601.290 or EN.601.340/440/640 or EN.601.421/621), or permission. Students may receive credit for only one of 601.425/625.

Prerequisite(s): Students may take EN.601.425 OR EN.601.625 for credit, but not both.

EN.601.626. Principles of Programming Languages. 3 Credits.

Same material as EN.601.426, for graduate students. Functional, object-oriented, and other language features are studied independent of a particular programming language. Students become familiar with these features by implementing them. Most of the implementations are in the form of small language interpreters. Some type checkers and a small compiler will also be written. The total amount of code written will not be overly large, as the emphasis is on concepts. The ML programming language is the implementation language used. Required course background: EN.601.226.

Prerequisite(s): Students may only receive credit for one of the following: EN.601.426 or EN.601.626.

Distribution Area: Engineering, Quantitative and Mathematical Sciences

EN.601.628. Compilers & Interpreters. 3 Credits.

Introduction to compiler design, including lexical analysis, parsing, syntax-directed translation, symbol tables, runtime environments, and code generation and optimization. Students are required to write a compiler as a course project. Required course background: intermediate programming, data structures and computer system fundamentals Recommended background: automata and computation theory

Prerequisite(s): Students may receive credit for only one of EN.601.428 or 601.628.

Distribution Area: Engineering

EN.601.629. Functional Programming in Software Engineering. 3 Credits.

How can we effectively use functional programming techniques to build real-world software? This course will primarily focus on using the OCaml programming language for this purpose. Topics covered include OCaml basics, modules, standard libraries, testing, quickcheck, build tools, functional data structures and efficiency analysis, monads, streams, and promises. Students will practice what they learn in lecture via functional programming assignments and a final project. Required course background in data structures (EN.601.226)

Prerequisite(s): Students can receive credit for EN.601.429 or EN.601.629, but not both.

Distribution Area: Engineering

EN.601.630. Combinatorics & Graph Theory in Computer Science. 3 Credits.

This is a graduate level course studying the applications of combinatorics and graph theory in computer science. We will start with some basic combinatorial techniques such as counting and pigeon hole principle, and then move to advanced techniques such as the probabilistic method, spectral graph theory and additive combinatorics. We shall see their applications in various areas in computer science, such as proving lower bounds in computational models, randomized algorithms, coding theory and pseudorandomness. Required Course Background: discrete math, probability theory and linear algebra.

Prerequisite(s): Students may receive credit for EN.601.430 or EN.601.630, but not both.

Distribution Area: Engineering, Quantitative and Mathematical Sciences

EN.601.631. Theory of Computation. 3 Credits.

This course covers the theoretical foundations of computer science. Topics included will be models of computation from automata to Turing machines, computability, complexity theory, randomized algorithms, inapproximability, interactive proof systems and probabilistically checkable proofs. Students may not take both 601.231 and 601.431/601.631, unless one is for an undergrad degree and the other for grad. Required Background: discrete math or permission; discrete probability theory recommended.

Prerequisite(s): Students can receive credit for only one of EN.601.431/EN.601.631

Distribution Area: Engineering, Quantitative and Mathematical Sciences

EN.601.633. Intro Algorithms. 3 Credits.

Same material as EN.601.433, for graduate students. This course concentrates on the design of algorithms and the rigorous analysis of their efficiency. topics include the basic definitions of algorithmic complexity (worst case, average case); basic tools such as dynamic programming, sorting, searching, and selection; advanced data structures and their applications (such as union-find); graph algorithms and searching techniques such as minimum spanning trees, depth-first search, shortest paths, design of online algorithms and competitive analysis. Required Background: data structures, discrete math, proof writing.

Prerequisite(s): Students may receive credit for only one of EN.600.363, EN.600.463, EN.601.433, EN.601.633

Distribution Area: Engineering, Quantitative and Mathematical Sciences

EN.601.634. Randomized and Big Data Algorithms. 3 Credits.

This course will cover fundamental methods of randomization for efficient algorithms and relevant techniques in probabilistic analysis. The first part of the course will discuss classical randomized algorithms, including the randomized algorithm for the min-cut problem, hashing techniques, tail inequalities, and the probabilistic method. The second part will delve into advanced topics such as coresets methods for clustering algorithms, the Johnson-Lindenstrauss lemma, and the applications of streaming and sketching algorithms.

Prerequisite(s): Students may receive credit for only one of EN.600.464, EN.600.664, EN.601.434, EN.601.634.

Distribution Area: Engineering

EN.601.635. Approximation Algorithms. 3 Credits.

Graduate version of EN.601.435. This course provides an introduction to approximation algorithms. Topics include vertex cover, TSP, Steiner trees, cuts, greedy approach, linear and semi-definite programming, primal-dual method, and randomization. Additional topics will be covered as time permits. There will be a final project. Recommended Background: EN.601.633 or equivalent. Students may receive credit for EN.601.435 or EN.601.635, but not both.

Prerequisite(s): Students can only receive credit for one of EN.601.435 or EN.601.635.

EN.601.636. Algorithmic Game Theory. 3 Credits.

Same material as EN.601.436, for graduate students. This course provides an introduction to algorithmic game theory: the study of games from the perspective of algorithms and theoretical computer science. There will be a particular focus on games that arise naturally from economic interactions involving computer systems (such as economic interactions between large-scale networks, online advertising markets, etc.), but there will also be broad coverage of games and mechanisms of all sorts. Topics covered will include a) complexity of computing equilibria and algorithms for doing so, b) (in)efficiency of equilibria, and c) algorithmic mechanism design.

Prerequisite(s): Students may receive credit for EN.601.436 or EN.601.636, but not both.

Distribution Area: Engineering, Quantitative and Mathematical Sciences

EN.601.638. Algorithmic Foundations of Differential Privacy. 3 Credits.

This course provides an introduction to differential privacy, with a focus on algorithmic aspects (rather than statistical or engineering aspects). Specific topics we will cover include: motivation for differential privacy, and different versions of differential privacy (pure, approximate, Renyi, and zero-concentrated in particular); basic mechanisms (Laplace, Gaussian, Discrete Gaussian, and Exponential); composition theorems; basic algorithmic techniques (sparse vector technique, private multiplicative weights, private selection); beyond global sensitivity: local sensitivity, propose-test-release, subsampling; differentially private graph algorithms; lower bounds. Required Course Background: 601.433/633 or permission. Students may receive credit for only one of 601.438/638.

Prerequisite(s): Students may earn credit for EN.601.438 or EN.601.638, but not both.

Distribution Area: Engineering, Quantitative and Mathematical Sciences

EN.601.640. Web Security. 3 Credits.

This course begins with reviewing basic knowledge of the World Wide Web, and then exploring the central defense concepts behind Web security, such as same-origin policy, cross-origin resource sharing, and browser sandboxing. It will cover the most popular Web vulnerabilities, such as cross-site scripting (XSS) and SQL injection, as well as how to attack and penetrate software with such vulnerabilities. Students will learn how to detect, respond, and recover from security incidents. Newly proposed research techniques will also be discussed. Required course background: data structures, computer system fundamentals and javascript/web development. Students may receive credit for only one of 601.340/440/640.

Prerequisite(s): Students may receive credit for only one of 601.340/440/640.

EN.601.641. Blockchains and Cryptocurrencies. 3 Credits.

Same as EN.601.441, for graduate students. This course will introduce students to cryptocurrencies and the main underlying technology of Blockchains. The course will start with the relevant background in cryptography and then proceed to cover the recent advances in the design and applications of blockchains. This course should primarily appeal to students who want to conduct research in this area or wish to build new applications on top of blockchains. It should also appeal to those who have a casual interest in this topic or are generally interested in cryptography. Students are expected to have mathematical maturity. Recommended Course Background: EN.601.226 AND (EN.553.310 OR EN.553.420)

Prerequisite(s): Students may receive credit for only one of EN.600.451 OR EN.601.441 OR EN.601.641

Distribution Area: Engineering

EN.601.642. Modern Cryptography. 3 Credits.

Same material as 601.442, for graduate students. Modern Cryptography includes seemingly paradoxical notions such as communicating privately without a shared secret, proving things without leaking knowledge, and computing on encrypted data. In this challenging but rewarding course we will start from the basics of private and public key cryptography and go all the way up to advanced notions such as zero-knowledge proofs, functional encryption and program obfuscation. The class will focus on rigorous proofs and require mathematical maturity. Required course background: Probability & Automata/Computation Theory

Prerequisite(s): Students may receive credit for only one of EN.601.442 OR EN.601.642.

Distribution Area: Engineering, Quantitative and Mathematical Sciences

EN.601.643. Security & Privacy in Computing. 3 Credits.

Same material as 601.443, for graduate students. Lecture topics will include computer security, network security, basic cryptography, system design methodology, and privacy. There will be a heavy work load, including written homework, programming assignments, exams and a comprehensive final. The class will also include a semester-long project that will be done in teams and will include a presentation by each group to the class. Required course background: C programming and computer system fundamentals.

Prerequisite(s): Students may receive credit for only one of EN.600.443, EN.601.443, EN.601.643.

Distribution Area: Engineering

EN.601.644. Medical Device Cybersecurity. 3 Credits.

In an increasingly connected healthcare landscape, medical devices have effectively become IT endpoints, often running general-purpose operating systems like Windows or Linux, incorporating cloud microservices, and integrating artificial intelligence to detect, prevent, and improve patient health outcomes. Protecting these devices from cyber threats is not just a technical challenge—it's a matter of patient safety. A security breach in medical devices like pacemakers or infusion pumps can have life-threatening consequences. National and international regulatory bodies, such as the FDA and EU National Competent Authorities (NCAs) and Medical Device Regulation (MDR), know the implications and have provided prescription and guidance emphasizing stringent cybersecurity measures throughout a device's lifecycle, from design and development to postmarket surveillance. The result is a heightened awareness of medical device security and its impact on healthcare delivery, requiring cybersecurity risk management. In particular, focusing on threat modeling, cybersecurity risk assessment, secure design, secure coding practices, vulnerability management and monitoring, software bill of materials, cybersecurity transparency, user labeling, penetration testing, and more. Recommended background: computing systems, operating systems, machine learning & AI. Students may receive credit for only one of 601.444/601.644.

Prerequisite(s): Students who have already taken, or are currently enrolled in EN.601.444, are not eligible to take EN.601.644.;EN.601.443 OR EN.601.643

EN.601.645. Practical Cryptographic Systems. 3 Credits.

Same material as 601.445, for graduate students. This semester-long course will teach systems and cryptographic design principles by example: by studying and identifying flaws in widely-deployed cryptographic products and protocols. Our focus will be on the techniques used in practical security systems, the mistakes that lead to failure, and the approaches that might have avoided the problem. We will place a particular emphasis on the techniques of provable security and the feasibility of reverse-engineering undocumented cryptographic systems.

Prerequisite(s): Students may receive credit for EN.600.454/EN.601.445 or EN.601.645, but not both.

Distribution Area: Engineering

EN.601.646. Sketching and Indexing for Sequences. 3 Credits.

Many of the world's largest and fastest-growing datasets are text, e.g. DNA sequencing data, web pages, logs and social media posts. Such datasets are useful only to the degree we can query, compare and analyze them. Here we discuss two powerful approaches in this area. We will cover sketching, which enables us to summarize very large texts in small structures that allow us to measure the sizes of sets and of their unions and intersections. This in turn allows us to measure similarity and find near neighbors. Second, we will discuss indexing — succinct and compressed indexes in particular — which enables us to efficiently search inside very long strings, especially in highly repetitive texts.

Prerequisite(s): Students may receive credit for EN.601.446 or EN.601.646, but not both.

EN.601.647. Computational Genomics: Sequences. 3 Credits.

Same material as 601.447, for graduate students. Your genome is the blueprint for the molecules in your body. It's also a string of letters (A, C, G and T) about 3 billion letters long. How does this string give rise to you? Your heart, your brain, your health? This, broadly speaking, is what genomics research is about. This course will familiarize you with a breadth of topics from the field of computational genomics. The emphasis is on current research problems, real-world genomics data, and efficient software implementations for analyzing data. Topics will include: string matching, sequence alignment and indexing, assembly, and sequence models. Course will involve significant programming projects. Required course background: Intermediate programming (C/C++ and Data Structures

Prerequisite(s): Students may receive credit for only one

EN.601.447/647/747

Distribution Area: Engineering

EN.601.649. Computational Genomics: Applied Comparative Genomics. 3 Credits.

The goal of this course is to study the leading computational and quantitative approaches for comparing and analyzing genomes starting from raw sequencing data. The course will focus on human genomics and human medical applications, but the techniques will be broadly applicable across the tree of life. The topics will include genome assembly & comparative genomics, variant identification & analysis, gene expression & regulation, personal genome analysis, and cancer genomics. The grading will be based on assignments, a midterm exam, class presentations, and a significant class project. Prerequisites: knowledge of the Unix operating system and programming expertise in a language such as R or Python.

Prerequisite(s): Students may receive credit for only one of EN.601.449/EN.601.649/EN.601.749.

EN.601.651. Introduction to Computational Immunogenomics. 3 Credits.

Immunology studies defensive mechanisms of living organisms against external threats. Computational immunogenomics is a new field of bioinformatics that develops and applies computational approaches to the study and interpretation of immunological data, seeking to answer questions about adaptive immune responses in humans and important animals. In this course, students will learn how to design, apply, and benchmark algorithms for solving immunogenomics problems. Required Course Background: Intermediate Programming & Data Structures. Students may receive credit for only one of EN.601.451, EN.601.651.

Prerequisite(s): Students may receive credit for only one of EN.601.451 OR EN.601.651.

EN.601.653. Applications of Augmented Reality. 3 Credits.

This course is designed to expand the student's augmented reality knowledge and introduce relevant topics necessary for developing more meaningful applications and conducting research in this field. The course addresses the fundamental concepts of visual perception and introduces non-visual augmented reality modalities, including auditory, tactile, gustatory, and olfactory applications. The following sessions discuss the importance of integrating user-centered design concepts to design meaningful augmented reality applications. A later module introduces the basic requirements to design and conduct user studies and guidelines on interpreting and evaluating the results from the studies. During the course, students conceptualize, design, implement and evaluate the performance of augmented reality solutions for their use in industrial applications, teaching and training, or healthcare settings. Homework in this course will relate to applying the theoretical methods used for designing, implementing, and evaluating augmented reality applications. Required course background: intermediate programming (C/C++), data structures, linear algebra; EN.601.654 preferred.

Prerequisite(s): Students may receive credit for only one of EN.601.453 or EN.601.653, but not both.;EN.601.454 OR EN.601.654

Distribution Area: Engineering

EN.601.654. Introduction to Augmented Reality. 3 Credits.

This course introduces students to the field of Augmented Reality. It reviews its basic definitions, principles, and applications. The course explains how fundamentals concepts of computer vision are applied for the development of Augmented Reality applications. It then focuses on describing the principal components and particular requirements to implement a solution using this technology. The course also discusses the main issues of calibration, tracking, multi-modal registration, advanced visualization, and display technologies. Homework in this course will relate to the mathematical methods used for calibration, tracking, and visualization in augmented reality. Required course background: intermediate programming (C/C++), data structures, linear algebra.

Prerequisite(s): Students may receive credit for only on EN.601.454/EN.601.654

Distribution Area: Engineering, Natural Sciences

EN.601.655. Computer Integrated Surgery I. 4 Credits.

Same material as 601.455, for graduate students. This course focuses on computer-based techniques, systems, and applications exploiting quantitative information from medical images and sensors to assist clinicians in all phases of treatment from diagnosis to preoperative planning, execution, and follow-up. It emphasizes the relationship between problem definition, computer-based technology, and clinical application and includes a number of guest lectures given by surgeons and other experts on requirements and opportunities in particular clinical areas. Required Course Background: data structures and linear algebra or permission. Recommended Course Background: intermediate programming in C/C++, EN.601.457, EN.601.461, image processing.

Prerequisite(s): Students may receive credit for only one of EN.601.455 or EN.601.655.

Distribution Area: Engineering

EN.601.656. Computer Integrated Surgery II. 3 Credits.

Same material as EN.601.456, for graduate students. This weekly lecture/seminar course addresses similar material to EN.601.655, but covers selected topics in greater depth. In addition to material covered in lectures/seminars by the instructor and other faculty, students are expected to read and provide critical analysis/presentations of selected papers in recitation sessions. Students taking this course are required to undertake and report on a significant term project under the supervision of the instructor and clinical end users. Typically, this project is an extension of the term project from EN.601.655, although it does not have to be. Grades are based both on the project and on classroom recitations. Students wishing to attend the weekly lectures as a 1-credit seminar should sign up for EN.601.356. Students may receive credit for only one of 601.456/496/656.

Prerequisite(s): Students may receive credit for only one of EN.600.446, EN.600.646, EN.601.456, EN.601.496, OR EN.601.656.;EN.600.445/ EN.601.455 OR EN.600.645/EN.601.655 OR permission of the instructor.

EN.601.657. Computer Graphics. 3 Credits.

Same material as 601.457, for graduate students. This course introduces computer graphics techniques and applications, including image processing, rendering, modeling and animation. Permission of instructor is required for students not satisfying a pre-requisite. No Audits. Required course background: EN.601.220 (C++), EN.601.226, linear algebra.

Prerequisite(s): Students may receive credit for only one of EN.601.457 OR EN.601.657.

Distribution Area: Engineering, Quantitative and Mathematical Sciences

EN.601.661. Computer Vision. 3 Credits.

This course provides an overview of fundamental methods in computer vision from a computational perspective. Methods studied include: camera systems and their modelling, computation of 3-D geometry from binocular stereo, motion, and photometric stereo, and object recognition, image segmentation, and activity analysis. Elements of machine learning and deep learning are also included. Required course background: Intro to Programming, Linear Algebra & prob/stats

Prerequisite(s): Students may receive credit for only one of EN.601.461, EN.601.661, OR EN.601.761.

Distribution Area: Engineering

EN.601.663. Algorithms for Sensor-Based Robotics. 3 Credits.

Same material as EN.601.463, for graduate students. This course surveys the development of robotic systems for navigating in an environment from an algorithmic perspective. It will cover basic kinematics, configuration space concepts, motion planning, and localization and mapping. It will describe these concepts in the context of the ROS software system, and will present examples relevant to mobile platforms, manipulation, robotics surgery, and human-machine systems. Recommended Course Background: EN.601.226, linear algebra, calculus, probability.

Prerequisite(s): Students may receive credit for only one of 601.463/663/763

Distribution Area: Engineering

EN.601.664. Artificial Intelligence. 3 Credits.

The course situates the study of Artificial Intelligence (AI) first in the broader context of Cognitive Science (i.e., human intelligence) and then treats in-depth principles and methods for reasoning, planning, and learning, including both conventional methods and recent deep learning approaches. The class is recommended for all scientists and engineers with a genuine curiosity about how to build an AI system (in particular, an intelligent agent) that can learn, reason about, and interact with the world and other agents. Students will be asked to complete both programming assignments and writing assignments. Required course background: data structures, linear algebra & prob/stat. Students can only receive credit for one of 601.464/664.

Prerequisite(s): Students may receive credit for only one of EN.601.464 OR EN.601.664.

Distribution Area: Engineering

EN.601.665. Natural Language Processing. 4 Credits.

Same material as 601.465, for graduate students. An in-depth introduction to core techniques for analyzing, transforming, and generating human language. The course spans linguistics, modeling, algorithms, and applications. (1) How should linguistic structure and meaning be represented (e.g., trees, morphemes, ?-terms, vectors)? (2) How can we formally model the legal structures and their probabilities (e.g., grammars, automata, features, log-linear models, recurrent neural nets, Transformers)? (3) What algorithms can estimate the parameters of these models (e.g., gradient descent, EM) and efficiently identify probable structures (e.g., dynamic programming, beam search)? (4) Finally, what kinds of systems can be built with these techniques and how are they constructed and evaluated in practice? Detailed assignments guide students through many details of implementing core NLP methods. The course proceeds from first principles, although prior exposure to AI, statistics, ML, or linguistics can be helpful. Prerequisite: Data Structures and basic familiarity with Python, partial derivatives, matrix multiplication and probabilities.

Prerequisite(s): Students may receive credit for only one of EN.601.465 OR EN.601.665.

Distribution Area: Engineering

EN.601.666. Information Retrieval and Web Agents. 3 Credits.

Same material as EN.601.466, for graduate students. An in-depth, hands-on study of current information retrieval techniques and their application to developing intelligent WWW agents. Topics include a comprehensive study of current document retrieval models, mail/news routing and filtering, document clustering, automatic indexing, query expansion, relevance feedback, user modeling, information visualization and usage pattern analysis. In addition, the course explores the range of additional language processing steps useful for template filling and information extraction from retrieved documents, focusing on recent, primarily statistical methods. The course concludes with a study of current issues in information retrieval and data mining on the World Wide Web. Topics include web robots, spiders, agents and search engines, exploring both their practical implementation and the economic and legal issues surrounding their use. [Applications] Recommended Course Background: EN.601.226

Prerequisite(s): Students can receive credit for EN.601.466 or EN.601.666, but not both

Distribution Area: Engineering

EN.601.667. Introduction to Human Language Technology. 3 Credits.

This course gives an overview of basic foundations and applications of human language technology, such as: morphological, syntactic, semantic, and pragmatic processing; machine learning; signal processing; speech recognition; speech synthesis; information retrieval; text classification; topic modelling; information extraction; knowledge representation; machine translation; dialog systems; etc. Required Background: EN.601.226 Data Structures; knowledge of Python recommended.

Prerequisite(s): Students can receive credit for only one of 601.467/667.

EN.601.668. Machine Translation. 3 Credits.

Same material as 601.468, for graduate students. Google translate can instantly translate between any pair of over fifty human languages (for instance, from French to English). How does it do that? Why does it make the errors that it does? And how can you build something better? Modern translation systems learn to translate by reading millions of words of already translated text, and this course will show you how they work. The course covers a diverse set of fundamental building blocks from linguistics, machine learning, algorithms, data structures, and formal language theory, along with their application to a real and difficult problem in artificial intelligence. Required course background: Data Structures and prob/stats

Prerequisite(s): Students may receive credit for only one of EN.601.468 OR EN.601.668.

Distribution Area: Engineering

EN.601.669. AI Safety, Alignment, & Governance. 3 Credits.

This course will focus on the alignment and governance challenges posed by advanced frontier/general purpose AI models: why these models may behave in ways that pose significant risk to human welfare and what technical and governance approaches might mitigate these risks. We'll begin the course studying general results from alignment and governance in human normative systems such as markets, politics, norms and laws. We'll pay special attention to risks arising from agentic AI. We'll then look at current technical and position papers in various topics in AI safety and alignment. Topics could include: RLHF, constitutional AI, red-teaming, safety evaluation methods, red lines, jail-breaking, prompt injection, over-optimization, and open-source debates. We'll conclude with discussion of regulatory frameworks such as regulatory markets, registration of frontier models, international governance organizations, registration of AI agents and legal personhood for AI agents. This is a paper-reading class.

Prerequisite(s): Students may only receive credit for EN.601.469 OR EN.601.669.

EN.601.670. Artificial Agents. 3 Credits.

This course covers a number of topics explored in introductory AI, such as knowledge representation, reasoning, and natural language understanding. Unlike introductory AI, we will pursue these topics based on the transformer neural architecture. We will motivate the material through interacting with agents in games: how to build models that understand user commands, how to generate responses back to a user, and how to reason about a synthetic environment to determine a course of action. Assignments will include programming, presentations on readings, and written summaries of readings.

EN.601.671. Natural Language Processing: Self-Supervised Models. 3 Credits.

The rise of massive self-supervised (pre-trained) models have transformed various data-driven fields such as natural language processing (NLP). In this course, students will gain a thorough introduction to self-supervised learning techniques for NLP applications. Through lectures, assignments, and a final project, students will learn the necessary skills to design, implement, and understand their own self-supervised neural network models, using the Pytorch framework. Required course background: data structures, linear algebra, probability, familiarity with Python/PyTorch, natural language processing or machine learning.

Prerequisite(s): Students may receive credit for EN.601.471 or EN.601.671, but not both.

Distribution Area: Engineering

EN.601.672. Natural Language Processing for Computational Social Science. 3 Credits.

[Alt. title: Analyzing Text as Data] Vastly available digitized text data has created new opportunities for understanding social phenomena. Relatedly, social issues like toxicity, discrimination, and propaganda frequently manifest in text, making text analyses critical for understanding and mitigating them. In this course, we will centrally explore: how can we use NLP as a tool for understanding society? Students will learn core and recent advances in text-analysis methodology, building from word-level metrics to embeddings and language models as well as incorporating statistical methods such as time series analyses and causal inference. Required Course Background: natural language processing and familiarity with Python/PyTorch. Students may receive credit for EN.601.472 or EN.601.672, but not both.

Prerequisite(s): Students can only receive credit for one of the following: EN.601.472 OR EN.601.672.

EN.601.673. Cognitive Artificial Intelligence. 3 Credits.

Humans, even young children, can learn, model, and reason about the world and other people in a fast, robust, and data efficient way. This course will discuss the principles of human cognition, how we can use machine learning and AI models to computationally capture these principles, and how these principles can help us build better AI. Topics will include (but are not limited to) Bayesian concept learning, probabilistic programming, intuitive physics, decision-making, Theory of Mind, pragmatics, and value alignment. Required Course Background: Prob/Stat & Linear Algebra & Computing; prior course in ML/ AI strongly recommended. Students may receive credit for only one of 601.473/601.673.

Prerequisite(s): Students who have taken EN.601.473 are not eligible to take EN.601.673.

EN.601.674. ML: Learning Theory. 3 Credits.

This machine learning course will provide a formal and in-depth coverage of topics in statistical and computational learning theory. We will revisit popular machine learning algorithms and understand their performance in terms of the size of the data (sample complexity), memory needed (space complexity), as well as the overall runtime (computational or iteration complexity). We will cover topics including PAC learning, uniform convergence, VC dimension, Rademacher complexity, algorithmic stability, kernel methods, online learning and reinforcement learning, as well as introduce students to current topics in large-scale machine-learning and randomized projections. General focus will be on combining methodology with theoretical and computational foundations.

EN.601.675. Machine Learning. 3 Credits.

Same material as 601.475, for graduate students. Machine learning is subfield of computer science and artificial intelligence, whose goal is to develop computational systems, methods, and algorithms that can learn from data to improve their performance. This course introduces the foundational concepts of modern Machine Learning, including core principles, popular algorithms and modeling platforms. This will include both supervised learning, which includes popular algorithms like SVMs, logistic regression, boosting and deep learning, as well as unsupervised learning frameworks, which include Expectation Maximization and graphical models. Homework assignments include a heavy programming components, requiring students to implement several machine learning algorithms in a common learning framework. Additionally, analytical homework questions will explore various machine learning concepts, building on the pre-requisites that include probability, linear algebra, multi-variate calculus and basic optimization. Students in the course will develop a learning system for a final project. Required course background: multivariable calculus, probability, linear algebra, intro to computing

Prerequisite(s): Students may receive credit for only one of EN.601.475 OR EN.601.675.

Distribution Area: Engineering

EN.601.677. Causal Inference. 3 Credits.

"Big data" is not necessarily "high quality data." Systematically missing records, unobserved confounders, and selection effects present in many datasets make it harder than ever to answer scientifically meaningful questions. This course will teach mathematical tools to help you reason about causes, effects, and bias sources in data with confidence. We will use graphical causal models, and potential outcomes to formalize what causal effects mean, describe how to express these effects as functions of observed data, and use regression model techniques to estimate them. We will consider techniques for handling missing values, structure learning algorithms for inferring causal directionality from data, and connections between causal inference and reinforcement learning. Pre-requisites: familiarity with the R programming language, multivariate calculus, basics of linear algebra and probability.

Prerequisite(s): Students may receive credit for only one of EN.601.477 OR EN.601.677.

Distribution Area: Engineering, Quantitative and Mathematical Sciences

EN.601.679. Machine Learning: Reinforcement Learning. 3 Credits.

Tremendous success of reinforcement learning (RL) in a variety of settings from AlphaGo to LLMs makes it a critical area to study. This course will study classical aspects of RL as well as its modern counterparts. Topics will include Markov Decision Processes, dynamic programming, model-based and model-free RL, temporal difference learning, Monte Carlo methods, multi-armed bandits, policy optimization and other methods. Required course background: machine learning, linear algebra and probability. Students may receive credit for at most one of 601.479/679.

Prerequisite(s): Students who have taken, or are currently enrolled in EN.601.479, are not eligible to take EN.601.679.

Distribution Area: Engineering

EN.601.682. Machine Learning: Deep Learning. 4 Credits.

Deep learning (DL) has emerged as a powerful tool for solving data-intensive learning problems such as supervised learning for classification or regression, dimensionality reduction, and control. As such, it has a broad range of applications including speech and text understanding, computer vision, medical imaging, and perception-based robotics. The goal of this course is to introduce the basic concepts of deep learning (DL). The course will include a brief introduction to the basic theoretical and methodological underpinnings of machine learning, commonly used architectures for DL, DL optimization methods, DL programming systems, and specialized applications to computer vision, speech understanding, and robotics. Students will be expected to solve several DL problems on standardized data sets, and will be given the opportunity to pursue team projects on topics of their choice. Required course background: Data Structures, Linear Algebra, Probability, Calc II required; Statistics, Machine Learning, Calc III, numerical optimization and Python strongly recommended.

Prerequisite(s): Students may receive credit for EN.601.482 or EN.601.682, but not both.

Distribution Area: Engineering

EN.601.684. Explainable AI Design & Human-AI Interaction. 3 Credits.

This is a design course. Increasing the trustworthiness of machine learning solutions has emerged as an important research area. One approach to trustworthy machine learning is explainable and/or interpretable machine learning, which attempts to reveal the working mechanisms of a machine learning system. However, other than on-task performance, explainability is not a property of machine learning algorithms, but an affordance: a relationship between explanation model and the target users in their context. Successful development of machine learning solutions that afford explainability thus requires understanding of techniques beyond pure machine learning. In this course, we will first review the basics of machine learning and human-centered design. Then, we will introduce several techniques to explain machine learning models and/or make them interpretable, and through hands-on sessions and case studies, will investigate how these techniques affect human-AI interaction. In addition to individual homework assignments, students will work in groups to design, justify, implement, and test an explainable machine learning algorithm for a problem of their choosing. Additional Recommended (601.454/654, 601.290, 601.490/690 or 601.491/691) and 601.477/677.

Prerequisite(s): Students may receive credit for EN.601.484 or EN.601.684, but not both.

EN.601.685. Probabilistic Models of the Visual Cortex. 3 Credits.

The course gives an introduction to computational models of the mammalian visual cortex. It covers topics in low-, mid-, and high-level vision. It briefly discusses the relevant evidence from anatomy, electrophysiology, imaging (e.g., fMRI), and psychophysics. It concentrates on mathematical modelling of these phenomena taking into account recent progress in probabilistic models of computer vision and developments in machine learning, such as deep networks. Programming experience (Python preferred).

Prerequisite(s): Students who have taken AS.050.375/AS.050.675 are not eligible to take EN.601.685.

Distribution Area: Quantitative and Mathematical Sciences

EN.601.686. Machine Learning: Artificial Intelligence System Design & Development. 3 Credits.

Advances in Artificial intelligence have opened new opportunities for developing systems to aid in numerous areas of society. In order for AI systems to succeed in making constructive and positive changes, we must consider their impact on everyday life. Specifically, AI system designers must evaluate the overall capabilities of the system, consider the resulting human-AI interactions, and ensure that the system behaves in a responsible and ethical manner. In this project-based course you will work in teams of 3-5 students to 1) Identify a need with high-impact implications on everyday life; 2) Articulate principles of Responsible AI relevant to the intended application, 3) Conceptualize and design an AI system targeting this need, and 4) Develop the AI system by refining a demo-able prototype based on feedback received during course presentations. Additionally, we will discuss potential ethical issues that can arise in AI and how to develop Responsible AI principles. Coursework will consist of writing assignments, project presentations, and a project demonstration. Required course background: (EN.601.475/675 or EN.601.464/664 or EN.601.482/682) and Python programming. Recommended: 601.290 or 601.454/654 or 601.490/690 or 601.491/691 (experience with human computer interface design).

Prerequisite(s): (EN.601.475 OR EN.601.675) OR (EN.601.464 OR EN.601.664) OR (EN.601.482 OR EN.601.682)

Distribution Area: Engineering

EN.601.687. Machine Learning: Coping with Non-Stationary Environments and Errors. 3 Credits.

This course teaches machine learning methods that 1) consider data distribution shift and 2) represent and quantify the model uncertainty in a principled way. The topics we will cover include machine learning techniques that deal with data distribution shift, including domain adaptation, domain generalization, and distributionally robust learning techniques, and various uncertainty quantification methods, including Bayesian methods, conformal prediction methods, and model calibration methods. We will introduce these topics in the context of building trustworthy machine learning solutions to safety-critical applications and socially-responsible applications. For example, a typical application is responsible decision-making under uncertainty in non-stationary environments. So we will also introduce concepts like fair machine learning and learning under safety constraints, and discuss how robust and uncertainty-aware learning techniques contribute to such more desired systems. Students will learn the state-of-the-art methods in lectures, test their understanding in homeworks, and apply these methods in a project. Required course background: 601.475/675 Machine Learning. Students may receive credit for only one of 601.487/687, and may not take this course after taking EN.601.787.

Prerequisite(s): Students who have taken or are enrolled in EN.601.787 OR EN.601.487 are not eligible to take EN.601.687.

EN.601.689. Human-in-the-Loop Machine Learning. 3 Credits.

Machine learning (ML) is being deployed in increasingly consequential tasks, such as healthcare and autonomous driving. For the foreseeable future, successfully deploying ML in such settings will require close collaboration and integration with humans, whether they be users, designers, engineers, policy-makers, etc. This course will look at how humans can be incorporated into the foundations of ML in a principled way. The course will be broken down into three parts: demonstration, collaboration, and oversight. Demonstration is about how machines can learn from 'observing' humans—such as learning to drive a car from data collected while humans drive. In this setting, the human is assumed to be strictly better than the machine and so the primary goal is to transmit the human's knowledge and abilities into the ML model. The second part, collaboration, is about when humans and models are near equals in performance but not in abilities. A relevant setting is AI-assisted healthcare: perhaps a human radiologist and ML model are good at diagnosing different kinds of diseases. Thus we will look at methodologies that allow machines to 'ask for help' when they are either unconfident in their own performance and/or think the human can better carry out the task. The course will close with the setting in which machines are strictly better at a task than humans are, but we still wish to monitor them to ensure safety and alignment with our goals (oversight). Assessment will be done with homework, quizzes, and a final project. Required course background: EN.601.475/675 Machine Learning or equivalent.

Prerequisite(s): Students who have taken or are enrolled in EN.601.489 are not eligible to take EN.601.689.

EN.601.690. Introduction to Human-Computer Interaction. 3 Credits.

This course is designed to introduce undergraduate and graduate students to design techniques and practices in human-computer interaction (HCI), the study of interactions between humans and computing systems. Students will learn design techniques and evaluation methods, as well as current practices and exploratory approaches, in HCI through lectures, readings, and assignments. Students will practice various design techniques and evaluation methods through hands-on projects focusing on different computing technologies and application domains. This course is intended for undergraduate and graduate students in Computer Science/Cognitive Science/Psychology. Interested students from different disciplines should contact the instructor before enrolling in this course. Recommended Background: Basic programming skills.

Prerequisite(s): Students can receive credit for either EN.601.490 or EN.601.690, but not both.

Distribution Area: Engineering

EN.601.691. Human-Robot Interaction. 3 Credits.

This course is designed to introduce graduate students to research methods and topics in human-robot interaction (HRI), an emerging research area focusing on the design and evaluation of interactions between humans and robotic technologies. Students will (1) learn design principles for building and research methods of evaluating interactive robot systems through lectures, readings, and assignments, (2) read and discuss relevant literature to gain sufficient knowledge of various research topics in HRI, and (3) work on a substantial project that integrates the principles, methods, and knowledge learned in this course. Pre-req: EN.601.220 and EN.601.226.

Prerequisite(s): Students may receive credit for EN.601.491 or EN.601.691.

EN.601.693. Accessible Computing. 3 Credits.

This course is designed to introduce students to the principles, challenges, and opportunities in designing computing systems that are accessible to people with diverse abilities. Students will learn about assistive technologies, inclusive design methodologies, and the incorporation of accessibility in computer applications through lectures, readings, and projects. Required Background: programming and knowledge in human-computer interaction.

Prerequisite(s): Students may only receive credit for EN.601.493 OR EN.601.693.

EN.601.713. Future Networks. 3 Credits.

Early networks were used for short message exchanges (Telegraph), and then the world moved to voice telephony. Today, the Internet's dominant traffic is entertainment video. More and more objects (IoT devices) are connected to the Internet for control and monitoring. With the need for enormous AI computations, new networks with gigantic capacity are being designed and built. These transformations require transferring large amounts of information, rapidly deploying new features, and simpler management. The course will start with a brief introduction to the past networks: telegraph and telephone networks. Then, it will move to today's Internet. Endpoints are not just humans but also objects and machines; the Internet is increasingly becoming a network of objects. The course will mostly focus on how these networks will evolve in the future. New applications such as autonomous driving require networking and computing to be embedded together. This feature is already beginning to be implemented in 5G and 6G networks; 6G will also allow networks to be used as sensors. New technologies such as mobile edge computing, software-defined networking (SDN), network slicing, digital twins, and named-data networking (NDN) enable these advances. Two timely topics – Web 3 and the application of machine learning to networking – have been added. V2X networks will be a strong focus. Students will be required to participate in discussions on this topic. Students will be asked to study new papers and do course projects, which should result in longer-term research projects. Recommended Course Background: A course in computer networks (e.g., EN.601.414/614 Computer Network Fundamentals).

EN.601.714. Advanced Computer Networks. 3 Credits.

This is a graduate-level course on computer networks. It provides a comprehensive overview on advanced topics in network protocols and networked systems. The course will cover both classic papers on Internet protocols and recent research results. It will examine a wide range of topics, e.g., routing, congestion control, network architectures, datacenter networks, network virtualization, software-defined networking, and programmable networks, with an emphasis on core networking concepts and principles. The course will include lectures, paper discussions, programming assignments and a research project. Recommended Course Background: One undergraduate course in computer networks (e.g., EN.601.414/614 Computer Network Fundamentals or the equivalent), or permission of the instructor. The course assignments and projects assume students to be comfortable with programming.

EN.601.715. Advanced Networks: Internet Measurement. 3 Credits.

This course will be an introduction to Internet measurement, and especially how it relates to security and policy. This course builds on the topics in EN 601.414/614, discussing vulnerabilities of internetworking protocols (BGP), the domain name system (DNS), and HTTPS certificate management. The goal of this course is to learn about current research, and get hands-on experience with real Internet measurement data. This data will help reveal the structure of the modern Internet, and the financial relationships that continue to shape it. Required Course Background: An undergraduate course in computer networks (e.g., EN.601.414/614 Computer Network Fundamentals or the equivalent), or permission of the instructor.

EN.601.716. Advanced Topics in Internet of Things. 3 Credits.

This course explores the convergence of computer networks, mobile computing, and embedded systems, with a specific focus on the Internet of Things (IoT). IoT represents a paradigm shift in computing, aiming to bridge the gap between the physical and digital worlds. Its development has opened up new possibilities, including mobile health, smart homes, industrial automation, and more. Throughout the course, students will delve into advanced topics such as IoT networking, mobile and edge computing, embedded machine learning, wireless sensing, human-computer interaction, and mobile health applications. To excel in this course, students are expected to engage in pre-class readings and in-class discussions, and complete a semester-long project. The focus of the course will be on training research philosophy and principles instead of papers' technical details. The course covers multiple disciplines and encourages interdisciplinary projects; students with diverse backgrounds such as computer science/engineering, electrical engineering, biomedical engineering or other related areas are welcomed. [Systems] Recommended Course Background: familiarity with computer system fundamentals, computer networks, signal processing, and mobile computing.

EN.601.717. Advanced Distributed Systems & Networks. 3 Credits.

The course explores the state of the art in distributed systems, networks and Internet research and practice, trying to see what it would take to push the envelope a step further. The course is conducted as a discussion group, where the professor and students brainstorm and pick interesting semester-long projects with high potential future impact. Example areas include robust scalable infrastructure (distributed datacenters, cloud networking, scada systems), real-time performance (remote surgery, trading systems), hybrid networks (mesh networks, 3-4G/Wifi/Bluetooth). Students should feel free to bring their own topics of interest and ideas. Recommended Course Background: a systems course (distributed systems, operating systems, computer networks, parallel programming) or permission of instructor.

EN.601.727. Machine Programming. 3 Credits.

Programs are the fundamental medium through which humans interact with computers. With the advent of large language models (LLMs), the automated synthesis of programs is rapidly transforming how we build software. Instead of manual code writing, we specify intent through examples, specifications, and natural language. This course explores both the foundations and frontiers of program synthesis, covering traditional symbolic techniques alongside LLM-driven approaches. Students will study a variety of synthesis paradigms, including example-based, type- and specification-guided, and interactive methods. We will examine how LLMs are applied to general-purpose programming tasks as well as to specialized domains such as theorem proving, program repair, planning, and verification. Throughout the course, students will gain exposure to a wide range of programming languages, from widely-used ones like Python and C, to emerging and domain-specific languages such as Rust, Lean, CodeQL, and PDDL. The course offers a research-oriented perspective combined with hands-on assignments and projects, providing students with both conceptual understanding and practical experience at the intersection of programming languages and machine learning. Required course background: Python proficiency and LLM familiarity.

EN.601.740. Language-based Security. 3 Credits.

This course will introduce Language-based Security, an emerging field in cyber security that leverages techniques from compilers and program analysis for security-related problems. Topics include but are not limited to: Control-flow and data-flow graphs, Program slicing, Code property graph (CPG), and Control-flow integrity. Students are expected to read new and classic papers in this area and discuss them in class. Recommended backgrounds are Operating Systems and preferably Compilers.

EN.601.742. Advanced Topics in Cryptography. 3 Credits.

This course will focus on advanced cryptographic topics with an emphasis on open research problems and student presentations.
Prerequisite(s): EN.601.442/EN.601.642 OR EN.601.445/EN.601.645 OR Permission of Instructor.

EN.601.743. Advanced Topics in Computer Security. 3 Credits.

Topics will vary from year to year, but will focus mainly on network perimeter protection, host-level protection, authentication technologies, intellectual property protection, formal analysis techniques, intrusion detection and similarly advanced subjects. Emphasis in this course is on understanding how security issues impact real systems, while maintaining an appreciation for grounding the work in fundamental science. Students will study and present various advanced research papers to the class. There will be homework assignments and a course project. A college level security or crypto course at Hopkins or any other school is required.

EN.601.760. FFT in Graphics & Vision. 3 Credits.

In this course, we will study the Fourier Transform from the perspective of representation theory. We will begin by considering the standard transform defined by the commutative group of rotations in 2D and translations in two- and three-dimensions, and will proceed to the Fourier Transform of the non-commutative group of 3D rotations. Subjects covered will include correlation of images, shape matching, computation of invariances, and symmetry detection. Recommended Course Background: AS.110.201 and comfort with mathematical derivations. No Audits.

EN.601.763. Advanced Topics in Robot Perception. 3 Credits.

The goal of this course is to explore machine learning and perception algorithms focused on robotic applications. Topics will include robot localization and mapping, pedestrian/obstacle detection/prediction, semantic segmentation, perception-based grasp planning, continual learning for perception algorithms and multimodal sensor fusion. This course will include introductions to the topics by the instructor followed by paper reading and discussions led by the students. In addition, this course will consist of an in-depth semester long project that will emphasize research skills including developing a hypothesis, conducting literature reviews, formulating the problem, defining, and conducting experiments and finally evaluating and reporting results. Required Course Background: Programming, Linear Algebra, Prob/Stat, Computer Vision and (Machine Learning or ML: Deep Learning).

Prerequisite(s): Students may only earn credit for one of the following: EN.600.336, EN.600.436/EN.601.463, EN.600.663, or EN.600.636/EN.601.763.

Distribution Area: Engineering, Natural Sciences

EN.601.764. Advanced NLP: Multilingual Methods. 3 Credits.

This is a project based course focusing on the design and implementation of systems that scale Natural Language Processing methods beyond English. The course will cover both multilingual and cross-lingual methods with an emphasis on zero-shot and few-shot approaches, as well as 'silver' dataset creation. Modules will include Cross-Lingual Information Extraction & Semantics, Cross-Language Information Retrieval, Multilingual Question Answering, Multilingual Structured Prediction, Multilingual Automatic Speech Recognition, as well as other non-English centric NLP methods. Students will be expected to work in small groups and pick from one of the modules to create a model based on state-of-the-art methods covered in the class. The course will be roughly two-thirds lecture based and one-third students presenting project updates periodically throughout the semester.

Prerequisite(s): EN.601.465 OR EN.601.665

Distribution Area: Engineering, Natural Sciences

EN.601.770. AI Ethics and Social Impact. 3 Credits.

AI is poised to have an enormous impact on society. What should that impact be and who should get to decide it? The goal of this course is to critically examine AI research and deployment pipelines, with in-depth examinations of how we need to understand social structures to understand impact. In application domains, we will examine questions like "who are key stakeholders?", "who is affected by this technology?" and "who benefits from this technology?". We will also conversely examine: how can AI help us learn about these domains, and can we build from this knowledge to design AI for "social good"? As a graduate-level course, topics will focus on current research including development and deployment of technologies like large language models and decision support tools, and students will conduct a final research project. Required Course Background: At least one graduate-level computer science course in Artificial Intelligence or Machine Learning (including NLP, Computer Vision, etc.), two preferred, or permission of the instructor.

EN.601.771. Advances in Self-Supervised Statistical Models. 3 Credits.

The rise of massive self-supervised (pre-trained) models has transformed various data-driven fields such as natural language processing, computer vision, robotics, and medical imaging. This advanced graduate course aims to provide a holistic view of the issues related to these models: We will start with the history of how we got here, and then delve into the latest success stories. We will then focus on the implications of these technologies: social harms, security risks, legal issues, and environmental impacts. The class ends with reflections on the future implications of this trajectory. Required Course Background: knowledge equivalent to EN.601.471/671 or EN.601.465/665; linear algebra and statistics.

Prerequisite(s): EN.601.471 OR EN.601.671 OR EN.601.465 OR EN.601.665

EN.601.773. Machine Social Intelligence. 3 Credits.

No other species possesses a social intelligence quite like that of humans. Our ability to understand one another's minds and actions, and to interact with one another in rich and complex ways, is the basis for much of our success, from governments to symphonies to the scientific enterprise. This course will discuss the principles of human social cognition, how we can use machine learning and AI models to computationally capture these principles, how these principles can help us build human-level machine social intelligence, and how social intelligence can enable the engineering of AI systems that can understand and interact with humans safely and productively in real-world settings. In this seminar course, we will read and discuss literature that cover diverse topics on social intelligence in humans and machines. These include (but are not limited to) social perception, Theory of Mind, multi-agent planning, multi-agent communication, social learning, human-AI teaming, moral judgment, and value alignment. Required Course Background: Linear Algebra, Probability and Statistics, and Calculus; 601.475/675 Machine Learning or EN.601.464/664 Artificial Intelligence or equivalent.

EN.601.774. Theory of Replicable Machine Learning. 3 Credits.

Replicability is vital to ensuring scientific conclusions are reliable, but failures of replicability have been a major issue in nearly all scientific areas of study, and machine learning is no exception. In this course, we will study replicability as a property of learning and other statistical algorithms, developing a theory of replicable learning. We will cover recent formalizations of replicability and their relationships to other common stability notions such as differential privacy and adaptive generalization. We will survey replicable algorithms for fundamental learning tasks, and discuss the limitations of replicable algorithms. If time permits, we will discuss replicability in other settings, such as reinforcement learning and clustering, or other useful and related stability notions such as list replicability and global stability. Required Course Background: EN.601.433/633 Intro Algorithms or instructor permission.

EN.601.778. Advanced Topics in Causal Inference. 3 Credits.

This course will cover advanced topics on all areas of causal inference, including learning causal effects, path-specific effects, and optimal policies from data featuring biases induced by missing data, confounders, selection, and measurement error, techniques for generalizing findings to different populations, complex probabilistic models relevant for causal inference applications, learning causal structure from data, and inference under interference and network effects. The course will feature a final project which would involve either an applied data analysis problem (with a causal inference flavor), a literature review, or theoretical work. Recommended Course Background: EN.601.447/677.

Prerequisite(s): EN.601.477 OR EN.601.677

EN.601.779. Machine Learning: Advanced Topics. 3 Credits.

This course will focus on recent advances in machine learning. Topics will vary from year to year. The course will be project focused and involve presenting and discussing recent research papers.

EN.601.783. Vision as Bayesian Inference. 3 Credits.

This is an advanced course on computer vision from a probabilistic and machine learning perspective. It covers techniques such as linear and non-linear filtering, geometry, energy function methods, markov random fields, conditional random fields, graphical models, probabilistic grammars, and deep neural networks. These are illustrated on a set of vision problems ranging from image segmentation, semantic segmentation, depth estimation, object recognition, object parsing, scene parsing, action recognition, and text captioning. Required course background: calculus, linear algebra (AS.110.201 or equiv.), probability and statistics (AS.550.311 or equiv.), and the ability to program in Python and C++.

EN.601.787. Advanced Machine Learning: Machine Learning for Trustworthy AI. 3 Credits.

This course teaches advanced machine learning methods for the design, implementation, and deployment of trustworthy AI systems. The topics we will cover include but are not limit to different types of robust learning methods, fair learning methods, safe learning methods, and research frontiers in transparency, interpretability, privacy, sustainability, AI safety and ethics. Students will learn the state-of-the-art methods in lectures, understand the recent advances by critiquing research articles, and apply/innovate new machine learning methods in an application. There will be homework assignments and a course project. Recommended course background: EN.601.475/675.

Distribution Area: Engineering

EN.601.788. Machine Learning for Healthcare. 3 Credits.

This course surveys the technical and practical challenges of applying machine learning in healthcare, focusing on two themes: The first theme will cover applications of machine learning to a wide range of healthcare data modalities (e.g., medical imaging, structured health records, etc). Beyond reviewing specific modeling approaches, we will focus on navigating pitfalls in model development and evaluation that arise in a healthcare context. The second theme will cover methodological approaches to developing safe and effective machine learning systems in healthcare, including topics such as (but not limited to) causality, fairness, and distribution shift. This course is designed for students who have a solid existing background in machine learning, and who are interested in both the technical and practical nuances of applying machine learning in healthcare. Grading will be done on the basis of homework assignments as well as a final project. Required course background: EN.601.475/675 Machine Learning or equivalent.

EN.601.790. Advanced Human-Computer Interaction: Research Methods. 3 Credits.

This course is specifically tailored for graduate students, especially PhD students, to provide them with a comprehensive review of Human-Computer Interaction (HCI) research. The course covers foundations and frontiers in HCI research. Core topics include interaction, social computing, and AI+HCI; breadth topics include collaboration, conversational interactions, ubiquitous and tangible computing, and accessibility. We will examine research methods, philosophies of research, and diverse ways of knowing to build foundational concepts and analytical skills for engaging in and understanding human-computer interaction research. Students will read research papers and methodological theory and engage in critical writing, group discussion, and oral presentations. Students will gain knowledge in learning and practicing commonly used methodologies in HCI, such as interview, field and lab experiment design, and qualitative and quantitative analysis methods. Required Course Background: 601.490/690 or permission.

EN.601.792. Advanced Topics in Conversational User Interfaces. 3 Credits.

As a critical component of human-computer interaction, conversational user interfaces (CUIs) have the potential to revolutionize the way we interact with technology. This course is designed for graduate students who want to gain a deeper understanding of CUIs and their real world applications. Throughout the course, students will explore cutting-edge research and methodologies for designing, implementing, and evaluating CUIs. Various forms of conversational interface will be covered, including chatbots, voice assistants, and multimodal dialogue systems. Coursework will include short open-ended assignments focused on applying methods learned in class, reading recent papers, and a course project. Required course background: EN.601.490/690 Intro HCI or permission.

Prerequisite(s): EN.601.490 OR EN.601.690

EN.601.794. Privacy Technology, Design, and Law. 3 Credits.

Privacy has long been considered as a fundamental human right. Emerging technologies such as social media, smart grid, Internet of Things, drones, and self-driving cars have raised heightened privacy issues. Recent developments of regulations such as the European Union's General Data Protection Regulation (GDPR) and California Consumer Privacy Act (CCPA) is also drawing increasing attention from technologists, policymakers, and the media. How to protect people's privacy is a key challenge of our time. This course provides an in-depth look into privacy, privacy laws and regulations, privacy-enhancing technologies and mechanisms, and privacy design. Privacy will be examined from historical, philosophical, cultural, legal, economic, behavioral, and technical perspectives. This course is designed primarily for graduate students who are interested in privacy and are from a wide range of disciplines such as information science, computer science and engineering, law, business, media studies, economics, politics, and psychology. Recommended course background: EN.601.443/643 Security & Privacy or equivalent.

EN.601.801. Computer Science Seminar. 1 Credit.

Seminar presentations by leading researchers across the field of computer science.

EN.601.803. Masters Research. 3 - 10 Credits.

Permission required. Independent research for masters or pre-dissertation PhD students.

EN.601.805. Graduate Independent Study. 1 - 3 Credits.

Permission required. Individual study in an area of mutual interest to a graduate student and a faculty member in the department.

EN.601.807. Teaching Practicum. 1 Credit.

PhD students will gain valuable teaching experience, working closely with their assigned faculty supervisor. Successful completion of this course fulfills the PhD teaching requirement.(grad students) Permission req'd.

EN.601.809. PhD Research. 3 - 20 Credits.

Independent research for PhD students.

EN.601.810. Diversity and Inclusion in Computer Science and Engineering. 1 Credit.

This reading seminar will focus on the question of diversity and inclusion in Computer Science (in particular) and engineering (in general). We aim to study the ways in which the curriculum, environment, and structure of computer science within academia perpetuates biases alienating female and minoritized students, and to explore possible approaches for diversifying our field. The seminar will meet on a weekly basis, readings will be assigned, and students will be expected to participate in the discussion.

EN.601.811. Future Faculty: Preparing a New Generation of PIs for the Academic Job Search. 1 Credit.

The goal of this seminar-style course is to prepare senior PhD students and postdocs in CS and robotics adjacent disciplines for the academic job market. At the end of the course sequence, it is expected that participants will 1) understand benefits and possible challenges of the academic career path, 2) be familiar with many aspects of the academic job market (such as timing, required documents, interview schedule, ...), 3) have completed a first draft of their application documents to be further refined with their respective advisors and mentors, 4) be prepared to tackle phone and on-campus interviews, and 5) have an appreciation for the essential tasks junior faculty must master quickly.

EN.601.817. Selected Topics in Systems Research. 1 Credit.

This course covers latest advances in the research of computer systems including operating systems, distributed system, mobile and cloud computing. Students will read and discuss recent research papers in top systems conferences. Each week, one student will present the paper and lead the discussion for the week. The focus topics covered in the papers vary semester to semester. Example topics include fault-tolerance, reliability, verification, energy efficiency, and virtualization.

EN.601.819. Selected Topics in Cloud Computing and Networked Systems. 1 Credit.

Participants will read and discuss seminal and recent foundational research on cloud and networked systems.

EN.601.826. Selected Topics in Programming Languages. 1 Credit.

This seminar course covers recent developments in the foundations of programming language design and implementation. Topics covered include type theory, process algebra, higher-order program analysis, and constraint systems. Students will be expected to present papers orally.

EN.601.831. CS Theory Seminar. 1 Credit.

Seminar series in theoretical computer science. Topics include algorithms, complexity theory, and related areas of TCS. Speakers will be a mix of internal and external researchers, mostly presenting recently published research papers.

EN.601.849. Selected Topics in Computational Immunogenomics. 1 Credit.

Immunology studies defensive mechanisms of living organisms against external threats. Computational immunogenomics is a new branch of bioinformatics that develops and applies computational approaches to the study and interpretation of immunological data, seeking to answer questions about human adaptive immune responses to various pathogens, including but not limited to flu, HIV, and SARS-CoV-2. In this course, students will attend lectures and present immunogenomics papers in a journal club format.

EN.601.856. Seminar: Medical Image Analysis. 1 Credit.

This weekly seminar will focus on research issues in medical image analysis, including imagesegmentation, registration, statistical modeling, and applications. It will also include selected topicsrelating to medical image acquisition, especially where they relate to analysis. The purpose of thecourse is to provide the participants with a thorough background in current research in these areas, as well as to promote greater awareness and interaction between multiple research groups withinthe University. The format of the course is informal. Students will read selected papers. All students will be assumed to have read these papers by the time the paper is scheduled for discussion. But individual students will be assigned on a rotating basis to lead the discussion on particular papers or sections of papers. Co-listed with En.520.746.

EN.601.857. Selected Topics in Computer Graphics. 1 Credit.

In this course we will review current research in computer graphics. We will meet for an hour once a week and one of the participants will lead the discussion for the week.

EN.601.862. Selected Topics in Medical Image Processing. 1 Credit.

This course will provide a background in medical imaging modalities and the unique aspects of image processing as it pertains to medical imaging. We will cover what an image is, how it is formed through six imaging modalities, and how images are typically stored, as well as background topics such as image metrics, quantification, filtering and transforms. More advanced topics will be discussed such as visualization, image enhancement, segmentation and registration. The final few weeks will introduce the topic of neural networks in image processing. Students will be expected to read and discuss publications, as well as complete an implementation project and report. Recommended course background: programming & linear algebra.

EN.601.864. Selected Topics in Multilingual Natural Language Processing. 1 Credit.

This is a weekly reading group focused on Natural Language Processing (NLP) outside of English. Whereas methods have gotten very strong in recent years on English NLP tasks, many methods fail on other languages due to both linguistic differences as well as lack of available annotated resources. This course will focus on Cross-Language Information Retrieval, Cross-Lingual Information Extraction, Multilingual Semantics, Massively Multilingual Language Modeling, and other non-English NLP sub-fields. Students will be expected to read, discuss, and present papers. Required course background: EN.601.465/665.

EN.601.865. Selected Topics in Natural Language Processing. 1 Credit.

A reading group exploring important current research in the field and potentially relevant material from related fields. In addition to reading and discussing each week's paper, enrolled students are expected to take turns selecting papers and leading the discussion.

Prerequisite(s): EN.601.465 OR EN.601.665.

EN.601.866. Selected Topics in Computational Semantics. 1 Credit.

This weekly reading group will review current research and survey articles on the topics of computational semantics, statistical machine translation, and natural language generation. Enrolled students will present papers and lead discussions.

EN.601.867. Selected Topics in Trustworthy & Responsible Natural Language Processing. 1 Credit.

This is a graduate student seminar aimed at introducing graduate students to the research areas of trustworthy and responsible NLP. This is primarily targeted at students with an NLP background and no or very little background in the course topics: bias, privacy, safety, misinformation, explainability, interpretability, robustness. Students will be expected to read, present and discuss papers each week. Required course background: NLP experience (eg. at least one of 601.665/667/668/671/765, etc.).

EN.601.868. Selected Topics in Machine Translation. 1 Credit.

Students in this course will review, present, and discuss current research in machine translation. Permission of instructor.

EN.601.888. Translate Health AI Seminar. 1 Credit.

Long title: Taking AI for healthcare from the GPU to the patient (THAI seminar).This seminar will cover various topics related to clinical translation of AI technologies in healthcare. Translation of AI for healthcare requires synergy across multiple domains including engineering, epidemiology and statistics, regulation, and clinical research methods. This series will include talks by faculty, invited speakers, and students on topics such as validation, model evaluation metrics, regulation, clinical evaluation including human machine interaction, studies on effectiveness and post-deployment real-world effects.