MS in Artificial Intelligence (AI)
Course Listings

PREPARATORY COURSES

You are encouraged to take preparatory courses that have been designed to help you fill limited background knowledge or skills fundamental to AI including programming or mathematical foundations. MS in AI students may take at most six graduate credits towards the degree of these preparatory courses:

CS 5007 Introduction to Programming Concepts, Data Structures, and Algorithms
This is an introductory graduate course teaching core computer science topics typically found in an undergraduate Computer Science curriculum, but at a graduate-level pace. It is primarily intended for students with little formal preparation in Computer Science to gain experience with fundamental Computer Science topics. After a review of programming concepts the focus of the course will be on data structures from the point of view of the operations performed upon the data and to apply analysis and design techniques to non-numeric algorithms that act on data structures. The data structures covered include lists, stacks, queues, trees, and graphs. Projects will focus on the writing of programs to appropriately integrate data structures and algorithms for a variety of applications. This course may not be used to satisfy degree requirements for a B.S., M.S., or Ph.D. degree in Computer Science or a minor in Computer Science. It may satisfy the requirements for other degree programs at the discretion of the program review committee for the particular degree.

Prerequisites: Experience with at least one high-level programming language such as obtained in an undergraduate programming course

CS 5008 Introduction to Systems and Network Programming
This course is focused on significant programming projects and provides an overview of the principles of computer networks and a general-purpose operating system. The course provides the student with an understanding of the basic components of an operating system, including processes, synchronization and memory management. The course exposes students to the Internet protocol suite networking layers while providing an introduction into topics such as wireless networking and Internet traffic considerations. The objective is to focus on an understanding of fundamental concepts of operating systems and computer network architecture from a design and performance perspective. Students will be expected to design and implement a variety of programming projects to gain an appreciation of the design of operating systems and network technologies. This course may not be used to satisfy degree requirements for a B.S., M.S., or Ph.D. degree in Computer Science or a minor in Computer Science. It may satisfy the requirements for other degree programs at the discretion of the program review committee for the particular degree.

Prerequisites: Experience with at least one high-level programming language such as obtained in CS 5007.

CS 5084 Introduction to Algorithms: Design and Analysis
This course is an introduction to the design, analysis and proofs of correctness of algorithms. Examples are drawn from algorithms for many areas. Analysis techniques include asymptotic worst case and average case, as well as amortized analysis. Average case analysis includes the development of a probability model. Techniques for proving lower bounds on complexity are discussed, along with NP-completeness. Note: students with a strong background in design and analysis of computer systems, at the level equal to a B.S. in computer science, should not take CS 5084 and should consider taking CS 504 or CS 584.

Prerequisites: An undergraduate knowledge of discrete mathematics and data structures.

DS 517/MA517 Mathematical Foundations for Data Science
The foci of this class are the essential statistics and linear algebra skills required for Data Science students. The class builds the foundation for theoretical and computational abilities of the students to analyze high dimensional data sets. Topics covered include Bayes' theorem, the central limit theorem, hypothesis testing, linear equations, linear transformations, matrix algebra, eigenvalues, and eigenvectors, and sampling techniques, including Bootstrap and Markov chain Monte Carlo. Students will use these techniques while engaging in hands-on projects with real data.

Prerequisites: Some knowledge of integral and differential calculus is recommended.

DS 501 Introduction to Data Science
Introduction to Data Science provides an overview of Data Science, covering a broad selection of key challenges in and methodologies for working with big data. Topics to be covered include data collection, integration, management, modeling, analysis, visualization, prediction, and informed decision making, as well as data security and data privacy. This introductory course is integrative across the core disciplines of Data Science, including databases, data warehousing, statistics, data mining, data visualization, high performance computing, cloud computing, and business intelligence. Professional skills, such as communication, presentation, and storytelling with data, will be fostered. Students will acquire a working knowledge of data science through hands-on projects and case studies in a variety of business, engineering, social sciences, or life sciences domains. Issues of ethics, leadership, and teamwork are highlighted.

Prerequisites: None beyond meeting the Data Science admission criteria.

DS 5006 Machine Learning for Engineering & Science Applications
This course surveys the application of data science (DS) and machine learning (ML) to problems arising in engineering and the sciences. While DS and ML have profoundly

affected domains such as image understanding and natural language processing, ML has seen comparatively less impact in chemistry, physics, chemical engineering, electrical engineering, and many other important application domains. Topics covered will include predictive modeling, feature engineering, and model assessment, with a particular focus on the small-data limit. We will analyze and apply algorithms with wide applicability in engineering and sciences including classic techniques such as multiple linear regression and random forests, and state-of-the-art techniques such as deep neural networks. Recommended Background:

Prerequisites: The intention is for the class to be accessible to a wide audience in disciplines outside of Computer Science and Data Sciences, though some basic background topics such as statistics or linear algebra, and the ability to learn Python programming at a basic level would be helpful.

MIS 587 Business Applications in Machine Learning
This course explores how Machine Learning (ML) and Artificial Intelligence (AI) is applied to solve business problems, to satisfy specific business needs, or to discover new opportunities for businesses. Applications of ML and AI are constantly evolving across many industries. This course utilizes existing AutoML solutions to address issues identified in business case studies (e.g. predicting hospital readmissions, loans likely to default, customer churn). The course covers the machine learning project life cycle starting with defining ML project objectives, acquiring, and exploring data, modeling using AutoML tools, interpretation of models and communication of outcomes, and implementation and deployment of predictive models in organizations.

CORE COURSES

MS in AI students must complete a five-course core by taking one course each in the five core MS-AI bins in AI, Ethics & AI, Machine Learning, Knowledge Representation & Reasoning, and Interaction & Action. Students may choose to take additional core courses, beyond the five required core courses, from below bins:

Artificial Intelligence (Choose One Course)

CS 534 Introduction to Artificial Intelligence
This course gives a broad survey of artificial intelligence. The course will cover methods from search, probabilistic reasoning, and learning, among other topics. Selected topics involving the applications of these tools are investigated. Such topics might include natural language understanding, scene understanding, game playing, and planning.

Fairness, Ethics, & Policy in AI (Choose One Course)

DS555/CS555 Responsible Artificial Intelligence

Artificial Intelligence (AI) algorithms have a significant impact on people's lives. In this course, we discuss social responsibility around data privacy, bias in data and decision making, policies as guardrails, fairness and transparency in the context of applying AI algorithms. Case studies considering societal challenges caused by AI technologies may include AI-based hiring recommendations stemming from societal biases present in training datasets, AI-empowered self-driving cars behaving in a dangerous manner when encountering atypical road conditions, digital health applications inadvertently revealing private patient information, or large language models like chat-GPT generating incorrect or harmful responses. This course also studies AI-based algorithmic solutions to some of these challenges. These include the design of robust machine learning algorithms with constraints to ensure fairness, privacy, and safety. Strategies for how to apply these methods to design safe and fair AI are introduced. Topics may include min-max optimization with applications to training machine learning models robust to adversarial attacks, stochastic methods for preserving privacy of sensitive data, and multi-agent machine learning models for reducing algorithmic bias and polarization in recommender systems.

Recommended Background: Pre-requisites: Machine Learning at the graduate level, undergraduate level (CS 4342), or), or equivalent knowledge.

SS560 Artificial Intelligence: Exploring Technology and Policy
In the rapidly evolving landscape of technology, artificial intelligence (AI) has emerged as a transformative force with significant implications for public policy. This course is designed to provide students with a comprehensive understanding of challenges and opportunities AI brings to society and attendant policy debates. Throughout the course, students will learn both knowledge and tools enabling them to critically analyze and contribute to policy formation, implementation, assessment of AI related polices. The course will prepare students to engage in ongoing policy and emergent policy debates that reflect AI's impact on society. Students will be able to make ethically informed decisions about the intersection of technology, policy, and society.

Prerequisites: None.

MIS520 Artificial Intelligence and its Ethical Application in Business
This course aims to provide the students with a comprehensive introduction to the recent developments in AI through the coverage of fundamental AI concepts and practical applications of these concepts in business. The course will allow students to understand AIs basic concepts and methods and apply AI-based techniques to solving practical business problems. Students will also experience how AI can transform businesses and gain an understanding of where AI technologies are heading within the next few years.

WR 513 Ethical Impact and Communication in Robotics and Artificial Intelligence Research

Engineers and other technologists are increasingly more aware of the ethical, legal, and social impacts of robotics and artificial intelligence. Some of them actively contribute to the creation and communication of new sets of ethical standards, such as the work done by IEEE's Global Initiative on Ethics of Autonomous and Intelligent Systems. What are the ethical principles that underpin these new standards? Since robots and AI systems are designed to work with or alongside humans, do people have a right to understand what autonomous systems are doing and why? How can roboticists and AI designers ensure that these systems are transparent and explainable? This course focuses on the communication of ethical and social impacts of scientific research and technology development. After learning about major debates in robot/AI/data ethics, students will cultivate skills to (1) conceptualize ethical inquiries in technology design and (2) articulate them in writing and other forms of scholarly communication. As part of this course, students will learn to apply the National Science Foundation's (NSF) broader impacts framework to their writing projects (dissertation, thesis, journal publication, grant application, etc.).

Machine Learning (Choose One)

CS 548 Knowledge Discovery and Data Mining
This course presents current research in Knowledge Discovery in Databases (KDD) dealing with data integration, mining, and interpretation of patterns in large collections of data. Topics include data warehousing and data preprocessing techniques; data mining techniques for classification, regression, clustering, deviation detection, and association analysis; and evaluation of patterns minded from data. Industrial and scientific applications are discussed.

Recommended Background: Background in artificial intelligence, databases, and statistics at the undergraduate level, or permission of the instructor. Proficiency in a high level programming language.

DS 502/MA543 Statistical Methods for Data Science
Statistical Methods for Data Science surveys the statistical methods most useful in data science applications. Topics covered include predictive modeling methods, including multiple linear regression, and time series, data dimension reduction, discrimination and classification methods, clustering methods, and committee methods. Students will implement these methods using statistical software.

Prerequisites: DS 517/ MA 517, Statistics at the level of MA 2611 and MA 2612 and linear algebra at the level of MA 2071.

CS 539 Machine Learning
The focus of this course is machine learning for knowledge-based systems. It will include reviews of work on similarity-based learning (induction), explanation-based learning, analogical and case-based reasoning and learning, and knowledge compilation. It will also

consider other approaches to automated knowledge acquisition as well as connectionist learning.

Prerequisites: CS 534 or equivalent, or permission of the instructor.

DS541/CS 541 Deep Learning
This course will offer a mathematical and practical perspective on artificial neural networks for machine learning. Students will learn about the most prominent network architectures including multilayer feedforward neural networks, convolutional neural networks (CNNs), auto-encoders, recurrent neural networks (RNNs), and generative-adversarial networks (GANs). This course will also teach students optimization and regularization techniques used to train them — such as back- propagation, stochastic gradient descent, dropout, pooling, and batch normalization. Connections to related machine learning techniques and algorithms, such as probabilistic graphical models, will be explored. In addition to understanding the mathematics behind deep learning, students will also engage in hands-on course projects. Students will have the opportunity to train neural networks for a wide range of applications, such as object detection, facial expression recognition, handwriting analysis, and natural language processing.

Prerequisites: Machine Learning (CS 539), and knowledge of Linear Algebra (such as MA 2071) and Algorithms (such as CS 2223).

CS586/DS504: Big Data Analytics
Big Data Analytics addresses the obstacle that innovation and discoveries are no longer hindered by the ability to collect data, but by the ability to summarize, analyze, and discover knowledge from the collected data in a scalable fashion. This course covers computational techniques and algorithms for analyzing and mining patterns in large-scale datasets. Techniques studied address data analysis issues related to data volume (scalable and distributed analysis), data velocity (high-speed data streams), data variety (complex, heterogeneous, or unstructured data), and data veracity (data uncertainty). Techniques include mining and machine learning techniques for complex data types, and scale-up and scale-out strategies that leverage big data infrastructures. Real-world applications using these techniques, for instance social media analysis and scientific data mining, are selectively discussed. Students are expected to engage in hands-on projects using one or more of these technologies.

Prerequisites: A beginning course in databases and a beginning course in data mining, or equivalent knowledge, and programming experience.

DS551/CS551 Reinforcement Learning
Reinforcement Learning is an area of machine learning concerned with how agents take actions in an environment with a goal of maximizing some notion of “cumulative reward”. The problem, due to its generality, is studied in many disciplines, and applied in many domains, including robotics and industrial automation, marketing, education and training,

health and medicine, text, speech, dialog systems, finance, among many others. In this course, we will cover topics including: Markov decision processes, reinforcement learning algorithms, value function approximation, actor-critics, policy gradient methods, representations for reinforcement learning (including deep learning), and inverse reinforcement learning. The course project(s) will require the implementation and application of many of the algorithms discussed in class.

ECE571 Machine Learning for Engineering Applications

This is an introductory course for engineering students to gain basic knowledge of machine learning and its applications. This course's objective is to learn machine learning theory and then apply it in engineering practice. A major emphasis of the course is to foster the capability of combining multiple machine learning techniques in complex problem solving, such as the detection of deepfake media. Topics include supervised learning, linear regression, kernel methods, support vector machine, neural networks, unsupervised learning, clustering, principal component analysis, deep learning with convolutional neural networks, and reinforcement learning. Students will develop software to implement machine learning and deep learning algorithms for practical engineering applications.

Prerequisites: Basic knowledge of probability and computer programming.

ECE557/CS557/DS557 Machine Learning for Cybersecurity

Machine Learning has proven immensely effective in a diverse set of applications. This trend has reached a new high with the application of Deep Learning virtually in any application domain. This course studies the applications of Machine Learning in the sub domain of Cybersecurity by introducing a plethora of case studies including anomaly detection in networks and computing, side-channel analysis, user authentication and biometrics etc. These case studies are discussed in detail in class, and further examples of potential applications of Machine Learning techniques including Deep Learning are outlined. The course has a strong hands-on component, i.e. students are given datasets of specific security applications and are required to perform simulations.

ECE556/CS556/DS556 On-Device Deep Learning

Deep Learning, a core of modern Artificial Intelligence, is rapidly expanding to resource-constrained devices, including smartphones, wearables, and intelligent embedded systems for improving response time, privacy, and reliability. This course focuses on bringing these powerful deep-learning applications from central data centers and large GPUs to distributed ubiquitous systems. On-Device Deep Learning is an interdisciplinary topic at the intersection of artificial intelligence and ubiquitous systems, dedicated to enabling computing on edge devices. This course includes a wide range of topics related to deep learning in resource constrained settings including pruning and sparsity, quantization, neural architecture search, knowledge distillation, on-device training and transfer learning, distributed training, gradient compression, federated learning, efficient data movement and accelerator design, dynamic network inference, and advanced compression and approximation techniques for enabling on-device deep neural network inference and training. This course provides a comprehensive foundation for cutting-edge “tinyML” expertise.

Recommended background: Prerequisites: The students should have an introductory undergraduate-level or graduate-level introductory background in machine learning and deep neural networks.

RBE577 Machine Learning for Robotics
This graduate-level course delves into the intersection of machine learning and robotics. The curriculum will explore the integration of contemporary learning techniques in robotic areas such as manipulation, navigation, planning, control, decision-making, and other pertinent challenges in robotics. Advanced deep learning techniques and their applications in robotics will be covered, including supervised learning (e.g., behavioral cloning, state prediction), reinforcement learning (e.g., actor-critic, visual foresight), and unsupervised/self-supervised methods (e.g., world model construction, learning forward dynamic models). In addition, the generalizability of these methods will be discussed, recent, and experimental studies will be conducted, examining the challenges of applying these techniques on physical systems.

Prerequisites: RBE 500 or equivalent
Recommended Background: RBE 501 and RBE 502.

Knowledge Representation & Reasoning (Choose One Course)

DS553/CS553 Machine Learning Development & Operations (ML OPS)
This course teaches students the computational skills required in the fields of Artificial Intelligence (AI) and Data Science. As data-driven decision-making and AI applications continue to transform industries, proficiency in programming and machine learning tools is important. In this course, you will develop a strong foundation in programming languages commonly used in AI and Data Science (such as Python). This course will cover the development, debugging, deployment, and subsequent monitoring phases of models in end-to-end pipelines core to machine learning systems. You will also familiarize yourself with popular libraries, frameworks and debugging on IDEs, such as PyCharm, PyTorch, scikit-learn, and/or pandas. Possible topics may include practice code development with a copilot as well as deployment of models on a cloud computing environment The student will engage in hands-on projects to practice their programming skills to solve real-world AI and Data Science problems.

Recommended Background: Prerequisites: Basic understanding of programming concepts, and concepts, and preferably some python knowledge.

CS542 Database Management Systems
An introduction to the theory and design of data-base management systems. Topics covered include internals of database management systems, fundamental concepts in

database theory, and database application design and development. In particular, logical design and conceptual modeling, physical database design strategies, relational data model and query languages, query optimization, transaction management and distributed databases. Typically there are hands-on assignments and/or a course project. Selected topics from the current database research literature may be touched upon as well.

Prerequisites: CS 5084 would be helpful

CS585/DS503 Big Data Management
Big Data Management deals with emerging applications in science and engineering disciplines that generate and collect data at unprecedented speed, scale, and complexity that need to be managed and analyzed efficiently. This course introduces the latest techniques and infrastructures developed for big data management including parallel and distributed database systems, map-reduce infrastructures, scalable platforms for complex data types, stream processing systems, and cloud-based computing. Query processing, optimization, access methods, storage layouts, and energy management techniques developed on these infrastructures will be covered. Students are expected to engage in hands-on projects using one or more of these technologies.

Prerequisites: A beginning course in databases at the level of CS 4432 or equivalent knowledge, and programming experience.

CS 509 Design of Software Systems
This course introduces students to a methodology and specific design techniques for team-based development of a software system. Against the backdrop of the software engineering life cycle, this course focuses on the object-oriented paradigm and its supporting processes and tools. Students will be exposed to industrial-accepted standards and tools, such as requirements elicitation, specification, modeling notations, design patterns, software architecture, integrated development environments and testing frameworks. Students will be expected to work together in teams in the complete specification, implementation and testing of a software application.

Prerequisites: Knowledge of a recursive high-level language and data structures. An undergraduate course in software engineering is desirable.

MIS 502 Data Management for Analytics
This course develops the skills business students need for handling data. It focuses on student skills in (1) cleaning and preparing data for analysis, (2) writing SQL queries to access and manipulate data, and (3) ethical uses of data and data privacy issues. It also covers the types of data typically found in organizations, e.g., employee, customer, product, marketing, operations, and financial data.

OIE 559 Advanced Prescriptive Analytics: From Data to Impact

This course provides an in-depth focus on prescriptive analytics, which involves the use of data, assumptions, and mathematical modeling of real-world decision problems to ascertain and recommend optimal courses of action. Starting from conceptualization of the problem, to using theory for translational modeling and techniques, to computational solving, and finally interpretation – likely in an iterative manner – students will gain knowledge of tools and practical skills in transforming real-world decision problems into actionable insights. Advanced topics in the prescriptive analytics domain will be covered, such as the use of integer variables to represent important logical constructs, using nonlinear functions to represent real-world decision aspects, the incorporation of stochasticity and uncertainty, and corresponding solution methods. Real-world problems will be selected from a variety of contexts that may include capacity management, data science, finance, healthcare, humanitarian operations, inventory management, production planning, routing, staffing, and supply chain. Students will complete an individual project that includes a report in the style of a technical report or research paper, as well as an oral presentation. Students may not receive credit for both OIE 4430 and OIE 559

Prerequisites: OIE 552, equivalent knowledge about optimization and linear programming, or consent of the instructor.

RBE 550 Robot Motion Planning
Motion planning is the study of algorithms that reason about the movement of physical or virtual entities. These algorithms can be used to generate sequences of motions for many kinds of robots, robot teams, animated characters, and even molecules. This course will cover the major topics of motion planning including (but not limited to) planning for manipulation with robot arms and hands, mobile robot path planning with non-holonomic constraints, multi-robot path planning, high-dimensional sampling-based planning, and planning on constraint manifolds. Students will implement motion planning algorithms in open-source frameworks, read recent literature in the field, and complete a project that draws on the course material. The PR2 robot will be available as a platform for class projects. Physical robot platforms will be available for class projects.

Prerequisites: Undergraduate Linear Algebra, experience with 3D geometry, and significant programming experience.

RBE 575 Safety and Guarantees in Autonomous Robotics
Robotic and AI systems have strong potential to directly impact our well-being, from self-driving cars to medical robots. Therefore, it is important to consider strong guarantees on the correctness and safety of their behavior. These guarantees ensure the robot will execute the desired behavior and will not execute undesired behavior. The course will define formal notions of system properties such as safety and liveness, explain how to model and analyze those properties in systems that make decisions and act on them, and understand the specific challenges related to making guarantees on embodied AI systems.

This course will cover many topics related to formal guarantees of safety and correctness in robotic and AI systems, including temporal logic-based planning, safe control via invariants and control barrier functions, neural net verification, closed-loop control with machine learning components, safe reinforcement learning, and other state-of-the-art topics at the intersection of safety, guarantees, AI, and robotics.

Prerequisites: RBE 500. Recommended Background: Machine Learning or Intro to AI.

RBE 511 Swarm Intelligence
This course will cover a wide range of topics in swarm intelligence, including mathematical, computational, and biological aspects. The course is organized in four parts. In the first part, the students will learn about complex systems and the basic concepts of self-organization, such as positive and negative feedback, symmetry breaking, and emergence. The second part concerns several types of network models, such as information cascades, epidemics, and voting. The instructor will illustrate a diverse collection of self-organized systems in nature, finance, and technology that concretize these concepts. The third part is dedicated to swarm robotics, and will cover common swarm algorithms for task allocation, collective motion, and collective decision-making. The fourth and final part covers optimization algorithms inspired by swarm intelligence, namely ant colony optimization and particles warm optimization. The course will blend theory and practice, challenging the students to learn by implementing the algorithms discussed in class through a final project in swarm robotics.

Recommended Background: C++/Python/Matlab programming; Linear algebra; Probability and statistics; Calculus.

Interaction & Action (Choose One Course)

DS 552/CS 552 Generative Artificial Intelligence
Generative Artificial Intelligence (Gen-AI) is a class of machine learning models that generate new data (text, images, faces, voice, artwork) that is near indistinguishable from the equivalent real data typically generated by humans. These models are trained based on realistic example data sets from the real world. This course covers the underlying fundamentals of generative models. It also introduces the design and modeling of some of the modern generative models: Variational Autoencoders (VAEs),7 Generative Adversarial Networks (GANs), Diffusion models, ChatGPT, Large Language Models, to name a few. Several applications will be discussed, ranging from image generation for engineering or science applications to the utilization of generated data for data augmentation in AI systems. Ethical concerns related to the danger of these generative technologies concerning issues from misinformation, bias, to data ownership are reviewed.

Recommended Background: Core artificial intelligence classes, such as machine learning and deep learning, or equivalent background is highly recommended.

DS 554/CS 554 Natural Language Processing
Natural Language Processing (NLP) is an interdisciplinary field at the intersection of artificial intelligence, linguistics, and computer science, dedicated to enabling computers to understand, interpret, and generate human language. NLP underpins advancements in human-computer interaction, information retrieval, sentiment analysis, chatbots, and a multitude of other applications. The course may cover a wide range of topics, including language modeling, sequence-to-sequence architectures, sentiment analysis, machine translation, and advanced techniques for natural language understanding and generation, providing a comprehensive foundation for NLP expertise.

Recommended Background: Programming skills at the level of CS 5007.

DS 547/CS 547 Information Retrieval
This course introduces the theory, design, and implementation of text-based and Web-based information retrieval systems. Students learn the key concepts and models relevant to information retrieval and natural language processing on large-scale corpus such as the Web and social systems. Topics include vector space model, crawling, indexing, web search, ranking, recommender systems, embedding and language model.

Prerequisites: statistical learning at the level of DS 502/MA 543 and programming skills at the level of CS 5007.

CS 549/RBE 549 Computer Vision
This course examines current issues in the computer implementation of visual perception. Topics include image formation, edge detection, segmentation, shape-from-shading, motion, stereo, texture analysis, pattern classification and object recognition. We will discuss various representations for visual information, including sketches and intrinsic images.

Prerequisites: CS 534, CS 543, CS 545, or the equivalent of one of these courses.

RBE 526/CS 526 Human-Robot Interaction
This course focuses on human-robot interaction and social robot learning, exploring the leading research, design principles and technical challenges we face in developing robots capable of operating in real-world human environments. The course will cover a range of multidisciplinary topics, including physical embodiment, mixed-initiative interaction, multi-modal interfaces, human-robot teamwork, learning algorithms, aspects of social cognition, and long-term interaction. These topics will be pursued through independent reading, class discussion, and a final project.

Prerequisites: Mature programming skills and at least undergraduate level knowledge of Artificial Intelligence, such as CS 4341. No hardware experience is required.) RBE 595 (Synergy of Human & Robot) and the RBE/CS 526 (Human-Robot Interaction) courses are equivalent. A student cannot take and get credit for both courses.

ECE 545/CS 545 Digital Image Processing
This course presents fundamental concepts of digital image processing and an introduction to machine vision. Image processing topics will include visual perception, image formation, imaging geometries, image transform theory and applications, enhancement, restoration, encoding and compression. Machine vision topics will include feature extraction and representation, stereo vision, model-based recognition, motion and image flow, and pattern recognition. Students will be required to complete programming assignments in a high-level language.

Prerequisites: Working knowledge of undergraduate level signal analysis and linear algebra; familiarity with probability theory is helpful but not necessary.

Capstone Project or MS Thesis

MS in AI students must complete either a three-credit capstone project experience **or** a nine-credit MS thesis **from the list below**.

For the capstone project, the MS-AI student can select one of the three capstone courses based on their primary interest and with approval of their MS-AI advisor and the instructor of the course.

- DS 598 Graduate Qualifying Project in Data Science (3 credits)
- CS 594/DS 594 Graduate Qualifying Project in Artificial Intelligence (3 credits)
- RBE 594 Capstone Project Experience in Robotics Engineering (3 credits)
- CS 599/DS 599/RBE 599 Master's Thesis (9 credits)

CS 594/DS 594 Graduate Qualifying Project in Artificial Intelligence (3 credits)
This three-credit graduate qualifying project, typically done in teams, provides a capstone experience in applying Artificial Intelligence skills to a real-world problem. It will be carried out in cooperation with an industrial sponsor and is approved and overseen by a core or collaborative faculty member in the Artificial Intelligence Program. This offering integrates theory and practice of Artificial Intelligence and includes the utilization of tools and techniques acquired in the Artificial Intelligence Program to a real-world problem. In addition to a written report, this project must be presented in a formal presentation to faculty of the AI program and sponsors. Professional development skills, such as communication, teamwork, leadership, and collaboration, will be practiced. This course is a degree requirement for the Master of Science in Artificial Intelligence (MS-AI) and may not be taken before completion of 21 credits in the program. Students outside the MS-AI program must get the instructor’s approval before.

Prerequisite: Completion of at least 24 credits of the AI degree, or consent of the instructor. With permission of the instructor, the GQP can be taken a 2nd time for a total of 6 credits.

CS 599/DS 599/RBE 599 Master's Thesis

The MS thesis in the Artificial Intelligence Program consists of a research or development project worth a minimum of 9 graduate credit hours. Students interested in research, and in particular those who consider pursuing a Ph.D. degree in a related area, are encouraged to select the M.S. thesis option. The student can sign up for MS thesis credits such as CS599, DS599, or RBE599, as long as a faculty affiliated with the MS-AI program serves as thesis advisor and the thesis topic relates to AI. Students must submit a thesis proposal for approval by the program by the end of the semester in which a student has registered for a third thesis credit and by the advisor. Proposals will be considered only at regularly scheduled program meetings. Students funded by a teaching assistantship, research assistantship or fellowship are expected to pursue the thesis option. The student then must satisfactorily complete a written thesis and present the results to the AI faculty in a public presentation.