

Computer Science (CS)

CS 6300. Principles of Artificial Intelligence. 3 Hours.

This course provides an understanding of Artificial Intelligence principles. Through lectures, hands-on projects, and real-world case studies, participants will gain practical knowledge of how AI systems are developed and applied across various industries. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Analyze problems using the agent-environment model to select appropriate artificial intelligence solution strategies. 2. Create artificial intelligence solutions to problems with the artificial intelligence pipeline, using algorithms and data structures of the field, including search, logic, probabilistic reasoning, and machine learning. 3. Apply unsupervised learning techniques, including dimensionality reduction and clustering, to identify patterns and refine features in data. 4. Perform statistical analysis of machine learning algorithms including linear and logistic regression. 5. Utilize statistical and probabilistic learning models such as Naive Bayes, Linear Discriminant Analysis, and Quadratic Discriminant Analysis for classification and regression. Prerequisites: Acceptance to the Graduate Certificate in Applied Artificial Intelligence and Machine Learning. FA.

CS 6310. Foundations of Machine Learning. 3 Hours.

This course discusses best practices in end-to-end machine learning, including data processing, model selection, validation, and production. This course also discusses various machine learning models including linear regression, ensemble methods, and neural networks. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Apply the core mathematics concepts required for machine learning, including multidimensional calculus, linear algebra, and probability, to analyze a machine learning algorithm. 2. Formulate a data driven question and implement an end-to-end machine learning pipeline to address that question. 3. Apply supervised learning models, including linear and logistic regression, decision trees, and support vector machines, to address regression and classification problems. 4. Utilize ensemble methods to improve model quality in machine learning pipelines. 5. Construct and utilize neural network models to address regression and classification problems, involving architecture selection, model optimization, visualization, and evaluation. Prerequisites: Acceptance to the Graduate Certificate in Applied Artificial Intelligence and Machine Learning. FA.

CS 6320. Foundations of Deep Learning. 3 Hours.

This course introduces foundational methods for computer vision and natural language processing. Students will understand the primary problems in both fields, such as segmentation, and apply deep learning frameworks to address these challenges. ***COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Construct convolutional neural network models for image classification. 2. Apply deep learning models for complex image tasks, including image segmentation. 3. Convert text to vectors, utilizing both word embedding and tokenizer frameworks. 4. Construct neural network models for natural language processing tasks such as semantic summarization, utilizing foundational models such as RNNs and LSTMs. 5. Fine tune and apply transformer models to advanced natural language processing tasks such as machine translation. Prerequisites: CS 6300 (Grade B- or higher); AND CS 6310 (Grade B- or higher). SP.

CS 6321. Natural Language Processing. 3 Hours.

This course covers core methods in natural language processing, including text preprocessing and representation, machine learning methods for language data, and pretrained language models (transformer-based). Students build, fine-tune, and evaluate models for tasks such as text classification, token-level labeling, summarization, and question answering. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Compare and justify appropriate text representation methods, including tokenization, embeddings, and contextual representations for specific NLP tasks. 2. Prepare and curate text datasets using normalization, segmentation, annotation, and data splitting strategies. 3. Implement and fine-tune neural NLP models, including recurrent and transformer-based approaches, for common NLP tasks. 4. Evaluate and interpret NLP systems using task-appropriate metrics and error analysis. 5. Analyze and address ethical issues in NLP application, including bias, privacy, and responsible deployment. Prerequisites: CS 6320 (Grade B- or higher) and admission to the MS in Artificial Intelligence program. FA.

CS 6322. Computer Vision. 3 Hours.

This course focuses on modern computer vision with deep learning, including visual data representation and preprocessing, deep learning methods for visual tasks, and vision architectures such as convolutional and transformer-based models. Students implement and evaluate models for tasks such as image classification, object detection, image segmentation, and visual representation learning. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Prepare image datasets with appropriate preprocessing, augmentation, and labeling strategies. 2. Construct convolutional neural network models for image classification and representation learning. 3. Implement and compare methods for object detection and image segmentation. 4. Evaluate and interpret computer vision systems using task-appropriate metrics and error analysis. 5. Apply responsible AI practices related to bias, privacy, and model robustness in computer vision. Prerequisites: CS 6320 (Grade B- or higher) and admission to the MS in Artificial Intelligence program. SP.

CS 6323. Advanced Deep Learning. 3 Hours.

This course examines advanced deep learning techniques, including modern architectures, optimization strategies, self-supervised learning, and generative modeling. Students implement, fine-tune, and evaluate state-of-the-art models across domains while considering scalability, robustness, and responsible use. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Compare advanced architectures such as transformers, diffusion models, and graph neural networks for appropriate tasks. 2. Apply optimization and regularization techniques to improve training stability and generalization. 3. Implement self-supervised or transfer learning workflows for data-scarce scenarios. 4. Evaluate model performance using rigorous experimental design and ablation studies. 5. Assess risks and limitations of advanced models, including bias, security, and compute constraints. Prerequisites: CS 6320 (Grade B- or higher) and admission to the MS in Artificial Intelligence program. SU.

CS 6330. Programming for Machine Learning in Life Sciences. 3 Hours.

Foundational course to computer programming for the life sciences, including data processing, visualization, and basic statistical analysis utilizing high level programming languages. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Create and modify modest computer programs using loops, conditionals, functions, and object oriented techniques. 2. Load, visualize, and preprocess data sets. 3. Construct programs for processing life science data using optimized data structures. 4. Model and analyze data using computational libraries. 5. Select an appropriate combination of programming modules to solve computing problems in the life sciences. Prerequisites: Acceptance to Graduate Certificate in Machine Learning for Life Sciences. FA.

CS 6331. Machine Learning for Life Sciences. 3 Hours.

Foundational course in end-to-end machine learning and frequently used machine learning models in the life sciences. Students will apply this knowledge in working with curated biological datasets. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Create machine learning solutions to problems in life sciences, using the machine learning pipeline, including data processing, model selection, model fine-tuning, and model validation. 2. Create and train various machine learning models, including: neural networks, regression models, and tree-based methods. 3. Select and justify use of a machine learning model for a specific problem. 4. Transform high dimensional datasets into lower dimensions for data analysis and exploration. 5. Test scientific hypotheses using statistical modeling. Corequisite: CS 6330. Prerequisites: Acceptance to the Graduate Certificate in Machine Learning for Life Sciences. FA.

CS 6341. Machine Learning for Drug Discovery. 3 Hours.

This course provides an in-depth exploration of the computational techniques and methodologies used in modern drug discovery. Students will learn practical applications of machine learning to drug discovery, focusing on drug development and design. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Describe the development and current-day applications of machine learning to drug discovery. 2. Apply standard chemistry, drug design and visualization libraries for molecular analysis. 3. Utilize machine learning models for problems in drug design, potentially including drug synergistic interactions, target identification, and ligand-receptor activation. 4. Create or repurpose potential drugs for treatment of specific diseases using machine learning models. Prerequisites: CS 6330 (B- or higher); AND CS 6331 (B- or higher). SP.

CS 6342. Machine Learning for Medical Imaging. 3 Hours.

This course explores the principles and techniques used in the processing and analysis of medical images. This course presents machine learning models for analysis of medical imaging data, including image processing, cell segmentation, image registration, and disease classification. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Perform basic medical imaging processing, including loading, saving, and transforming images. 2. Construct and apply convolutional neural network architectures and related techniques for classification of images. 3. Create machine learning architectures for image segmentation and registration of medical images. 4. Construct machine learning models for identification and diagnosis of disease from medical images. Prerequisites: CS 6330 (B- or higher); AND CS 6331 (B- or higher). SP.

CS 6343. Machine Learning for Genomics, Transcriptomics and Proteomics. 3 Hours.

This course will assess the relationships among sequence, structure, and function in complex biological systems and networks. This course covers the application of computational algorithms in genomic, transcriptomic, and proteomic data analysis. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Employ specialized software and libraries to load, process, and analyze genomic, transcriptomic and proteomic data. 2. Understand and implement genome annotation algorithms to predict genes, regulatory elements, or other functional regions in genomic sequences. 3. Create and apply machine learning models for precision medicine focusing on identification of potential disease-associated biomarkers. 4. Implement computational algorithms for protein structure, function, and interaction prediction. Prerequisites: CS 6330 (B- or higher); AND CS 6331 (B- or higher). SP.

CS 6349R. Special Topics in Machine Learning for Life Sciences. 1-3 Hours.

This course allows the student to explore the application of machine learning principles and techniques to life science problems not addressed in existing courses. The student will collect and process data, select and train learning models, and validate the results. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Perform data collection and preprocessing on data appropriate to the project. 2. Construct and apply appropriate machine learning models for the project. 3. Extract results from the data appropriate to the project.

CS 6350. Artificial Intelligence and Machine Learning Project 1. 1-3 Hours.

This course will guide students through an introductory machine learning project, enabling them to learn and practice required skills for problem formulation, data analysis, presentation of results, and model deployment. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Formulate a clear and actionable research question that can be addressed using machine learning techniques. 2. Collect relevant datasets and perform necessary preprocessing and feature engineering to prepare data for analysis. 3. Conduct exploratory data analysis and utilize data visualization to inform model selection and development. 4. Validate and fine-tune machine learning models using tools such as hyperparameter tuning and standard tools for performance evaluation. 5. Communicate findings effectively through reports and presentations and understand the basics of model deployment in real-world applications. Prerequisites: Acceptance to the Graduate Certificate in Applied Artificial Intelligence and Machine Learning and/or the Masters in Artificial Intelligence program. FA, SP.

CS 6351. Artificial Intelligence and Machine Learning Project 2. 1-3 Hours.

This course will guide students through a deep learning project, focusing particularly on either computer vision or natural language processing, enabling them to learn and practice required skills for problem formulation, data analysis, presentation of results, and model deployment. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Formulate a clear and actionable research question that can be addressed using deep learning techniques. 2. Collect relevant datasets and perform necessary preprocessing and feature engineering to prepare data for analysis. 3. Conduct exploratory data analysis and utilize data visualization to inform model selection and development. 4. Validate and fine-tune deep learning models using tools such as hyperparameter tuning and architecture search, as well as standard tools for performance evaluation. 5. Communicate findings effectively through reports and presentations and understand model deployment in real-world applications. Prerequisites: CS 6350 (Grade B- or higher). SP.

CS 6352. Artificial Intelligence and Machine Learning Project 3. 1-3 Hours.

This project-based course advances students through an applied AI project with increasing scope and technical depth, including machine learning or agentic AI systems. Students develop a proposal, implement baseline and advanced approaches, evaluate results, and document outcomes in professional technical reports following program guidelines. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Define a feasible AI project (machine learning or agentic) with clear inputs, outputs, and evaluation metrics. 2. Implement baseline and advanced approaches and justify model or agent design choices. 3. Evaluate systems with appropriate quantitative metrics and error analysis. 4. Produce professional project documentation that communicates methods, results, and limitations. 5. Apply reproducible workflows, including data management, experiment tracking, and version control. Prerequisites: CS 6351 (Grade B- or higher) and admission to the MS in Artificial Intelligence Program. FA, SP.

CS 6353. Artificial Intelligence and Machine Learning Project 4 (Capstone). 1-3 Hours.

This capstone project course completes the AI project sequence. Students design, implement, and evaluate a substantial applied AI project, including machine learning or agentic AI systems, culminating in a professional report and presentation that demonstrate technical depth, rigor, and reproducibility. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Formulate and execute a capstone-level AI project (machine learning or agentic) with a clear scope and measurable success criteria. 2. Implement and refine models or agents through systematic experimentation and evaluation. 3. Synthesize results into a cohesive technical report and presentation for a professional audience. 4. Critically assess limitations, ethical considerations, and future work. 5. Deliver a reproducible project artifact including code, data documentation, and evaluation evidence. Prerequisites: CS 6352 (Grade B- or higher) and admission to the MS in Artificial Intelligence program. SP, SU.

CS 6359R. Artificial Intelligence and Machine Learning Independent Project. 1-3 Hours.

In this course students will design and complete a project using AI and ML techniques. The project will address a real-world problem. Students will improve their skills in problem formulation, data analysis, presentation of results, and model deployment. ****COURSE LEARNING OUTCOMES (CLOs)**** At the successful conclusion of this course students will: 1. Formulate a clear and actionable research question that can be addressed using machine learning or artificial intelligence techniques. 2. Collect relevant datasets and perform necessary preprocessing and feature engineering to prepare data for analysis. 3. Conduct exploratory data analysis and utilize data visualization to inform model selection and development. 4. Validate and fine-tune learning models using tools such as hyperparameter tuning and architecture search, and standard tools for performance evaluation. 5. Communicate findings effectively through reports and presentations and understand model deployment in real-world applications. Prerequisites: CS 6350(B- or higher).