ARTIFICIAL INTELLIGENCE

The part-time Artificial Intelligence program will educate and train practicing scientists and engineers to be able to carry out engineering and scientifically oriented research and development using their artificial intelligence knowledge and skills.

The rigorous curriculum will provide engineers and computer scientists with a working knowledge of the theoretical concepts in artificial intelligence and will also provide the students with the knowledge and skills to apply both current and future theoretical concepts to real systems and processes. The course content will be based on the foundational content embodied in the current computer science courses modified to provide relevant examples in the artificial intelligence setting.

Program Committee

Barton Paulhamus, Program Chair
Principal Professional Staff
JHU Applied Physics Laboratory

Anthony N. Johnson, Program Manager
Senior Professional Staff
JHU Applied Physics Laboratory

Eleanor Boyle Chlan
Senior Professional Staff (retired)
JHU Applied Physics Laboratory

Diana Gehlaus
Research Fellow
Center for Security and Emerging Technology

Lanier Watkins, Chair Computer Science and Cybersecurity
Principal Professional Staff
JHU Applied Physics Laboratory

Programs

- Artificial Intelligence, Graduate Certificate (https://e-catalogue.jhu.edu/engineering/engineering-professionals/artificial-intelligence/graduate-certificate/)
- Artificial Intelligence, Master of Science (https://e-catalogue.jhu.edu/engineering/engineering-professionals/artificial-intelligence/master-of-science/)

Courses

EN.705.601. Applied Machine Learning. 3 Credits.
Machine Learning (ML) is the art of solving a computation problem using a computer without an explicit program. ML is now so pervasive that various ML applications such as image recognition, stock trading, email spam detection, product recommendation, medical diagnosis, predictive maintenance, cybersecurity, etc. are constantly used by organizations around us, sometimes without our awareness.In this course, we will rigorously apply machine learning techniques to real-world data to solve real-world problems. We will briefly study the underlying principles of diverse machine learning approaches such as anomaly detection, ensemble learning, deep learning with a neural network, etc. The main focus will be applying tool libraries from the Python-based Anaconda and Java-based Weka data science platforms to datasets from online resources such as Kaggle, UCI KDD, open source repositories, etc. We will also use Jupyter notebooks to present and demonstrate several machine learning pipelines.
Prerequisite(s): EN.705.621 Introduction to Algorithms OR EN.605.621 Foundations of Algorithms OR EN.685.621 Algorithms for Data Science

EN.705.603. Creating AI-Enabled Systems. 3 Credits.
Achieving the full capability of AI requires a system perspective, extending beyond the models, to effectively leverage algorithms, data, and computing power. Creating AI-enabled systems includes thoughtful consideration of an operational decomposition for AI solutions, engineering data for algorithm development, and deployment strategies. The objective of this course is to bring a system perspective to creating AI-enabled systems. The course will explore the full-lifecycle of creating AI-enabled systems starting with problem decomposition and addressing data, development, design, diagnostic, and deployment phases. Each module will either introduce a domain in Machine Learning (Tabular, Computer Vision, Natural Language Processing, and Physical Systems) or delve into the end-to-end development of a specific AI system. Students will be exposed to the common technologies and resources practitioners use to develop these systems.
Prerequisite(s): Algorithms and Machine Learning

EN.705.604. Optimizing and Deploying Scalable AI Systems. 3 Credits.
Build production AI systems through four critical skills: (1) Automating model development via data/model pipelines (2) Optimizing for performance and computing resource usage through experimentation, incorporating explanations, and leveraging shared/existing models (3) Form a complete system through the integration of supporting technologies such as high speed messaging, real-time APIs, parallel processing, data storage, and system management (4) Packaging into a deployed system via containerizing and orchestrating including post-deployment adaptation across heterogeneous computing fabrics. The student acquires the skills through developing a realistic, hands-on, collaborative, and incremental project that produces optimized models integrated into an operational system deployed across hybrid computer environments.
Prerequisite(s): Working knowledge of Python, and Machine Learning Model development from course EN.705.603.

EN.705.605. Introduction to Generative AI. 3 Credits.
This introductory course on generative artificial intelligence (AI) offers a comprehensive overview of the foundational principles and techniques that enable machines to generate complex outputs, such as text, images, and music. Students will explore the history and evolution of generative AI, including landmark models and milestones that have shaped the field. The course will first review classical approaches such as expert systems, genetic algorithms, Markov models, and constraint satisfaction problems. This will be followed by key algorithms and models in generative AI, including but not limited to neural networks, Autoencoders, generative adversarial networks (GANs), and transformers. Ethical considerations and societal impacts of generative AI will be a critical component of the curriculum, encouraging students to think critically about issues such as bias, fairness, and privacy. Assessment will include project work, where students will demonstrate their ability to develop a generative AI application from concept to implementation. By the end of the course, students will have a solid understanding of generative AI, equipped with the skills to innovate and apply these technologies in diverse and ethically responsible ways.

EN.705.612. Values and Ethics in Artificial Intelligence. 3 Credits.
Modern artificial intelligence, and the related area of autonomous systems are becoming so powerful that they raise new ethical issues. This course will prepare professional engineers and developers to thoughtfully engage with the moral, ethical, and cultural aspect of these emerging technology. Topics include: safety considerations for autonomous vehicles, algorithm bias, AI explainability, data privacy, ethical considerations of 'deep fakes', ethics of artificial life, values advocacy within organizations, technological unemployment, and far-future considerations related to AI safety.

EN.705.615. Artificial Intelligence for Leaders. 3 Credits.
This course is designed for leaders tasked with spearheading artificial intelligence (AI) efforts within their organizations. As AI technologies such as machine learning, deep learning, symbolic AI and generative AI reshape the landscape of industry and governance, understanding how to effectively integrate these tools into business strategies becomes paramount. This course offers an in-depth exploration of the critical components of AI, including data acquisition and analysis, algorithm development, the deployment of resources, labor considerations and the management of at-scale AI projects. Participants will gain a robust understanding of the foundational and advanced concepts of AI, including the workings of machine learning models, the revolutionary capabilities of transformers and large language models (LLMs), the innovative potential of generative AI, and risk mitigation with symbolic AI. The curriculum emphasizes not only the technical aspects but also the management and ethical considerations, such as bias mitigation and the development of responsible AI frameworks, ensuring leaders can make informed, ethical decisions in deploying AI technologies.

EN.705.621. Introduction to 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.

EN.705.640. Cognitive and Behavioral Foundations for Artificial Intelligence. 3 Credits.
As a result of greater computing power and Big Data, artificial intelligence (AI) is rapidlyimproving for well-defined tasks and narrow intelligence. Moreover, it has entered all industriesin a myriad of ways. But will AI ever have human-like general intelligence? What does humanlikegeneral intelligence even mean? Why should we even care? This course is designed toanswer these complex questions by giving students working knowledge of the underlyingprinciples and mechanisms of human behavior and cognition, and how they may be applied tosolving current and rising industry challenges. Key topics to be addressed will include vision,audition, language, learning, emotion and social cognition, creativity, and consciousness. Students will apply learned topics to a final group research project on the topic of their choice.

EN.705.643. Deep Learning Developments with PyTorch. 3 Credits.
PyTorch is a machine learning framework based on the Torch library. Its flexibility and user-friendliness have accumulated a massive user base in both industry and academia. Most modern research code is written in PyTorch. In this course, we will provide a step-by-step comprehensive coverage of modern applications in PyTorch. The course topics can be broadly categorized into three popular applications: computer vision, natural language processing, and reinforcement learning. We will study the experimental details of using PyTorch for a wide variety of tasks such as image/video classification, object detection, semantic segmentation, text classification, sequence-to-sequence translation, visual question answering, and DQN. In terms of modern deep learning architectures, we will cover 2D/3D convolutional neural networks, recurrent neural networks, long-short term memory, transformers, and encoder-decoder networks. Students will be technically prepared for more advanced courses in different application after taking this course.

EN.705.651. Large Language Models: Theory and Practice. 3 Credits.
An apparently new breed of neural network -- the large language model (LLM) -- figures increasingly in today's news: ChatpGPT and Microsoft's new chatbot-like Bing Chat interface seem to garner headlines on the daily. This course constitutes a thorough introduction to this technology, tracing the historical threads in computational linguistics and language modeling that led to it, and exploring the design patterns that underpin its application in modern AI systems. In between, students will learn about language modeling, the attention mechanism, prompt and instruction tuning, composability, quantization, low-rank adaptation, and the wealth of software and hardware optimizations that enable LLMs to be used at scale and with acceptable latencies.

EN.705.741. Reinforcement Learning. 3 Credits.
This course will focus on both the theoretical and the practical aspects of designing, training, and testing reinforcement learning systems. The course begins with an examination of Markov decision processes (MDPs), which provide a sound mathematical basis for modeling and solving complex sequential decision problems. The more traditional analytical method for solving MDPs, dynamic programming, will be reviewed. We will then examine the major reinforcement learning approaches, such as Monte Carlo methods, temporal difference methods, policy gradient methods, and deep learning methods, comparing them as appropriate to dynamic programming techniques. Fundamental issues and limitations on the performance of reinforcement learning algorithms (e.g., the credit assignment problem, the exploration / exploitation tradeoff, on-policy learning versus off-policy learning, partial observability, and algorithm convergence properties) will be examined for each approach. Weekly exercises and discussion topics will reinforce and expand on the classroom material. In addition, students will gain practical experience during a semester-long project by programming, training, and testing various reinforcement learning algorithms.
Prerequisite(s): EN.625.638/EN.605.647 - Neural Networks or experience programming artificial neural networks in a high-level language.

EN.705.742. Advanced Applied Machine Learning. 3 Credits.
Machine learning is a subset of artificial Intelligence to build and utilize data models based on sound analytical algorithms. Still, it takes more than just applying a set of algorithms to datasets or experiment a list of toolbox library to successfully build effective machine learning subsystems in an AI system. In this course, we will study a variety of advanced topics involving solutions and novel techniques to various machine learning problems. Starting from Machine Learning Operations, these topics include model analysis such as Recommender Systems, Hyperparameter Optimization, Transfer Learning, and Explainable AI. Moreover, we will study and implement Neural Network machine learning algorithms such as Generative Adversarial Networks, Recurrent Neural Networks, Transformers, and Graph Neural Networks. The course will keep a balance between the theoretical and mathematical specifications of an algorithm and the actual engineering of an algorithm. In addition, we will apply these methods and models, such as GPT, to a variety of real-world problems in realistic course assignments. The course will also keep a research thread with discussions about recent developments, and emerging technologies in the current literature. Students will be expected to write a research paper throughout the course.
Prerequisite(s): EN.705.601 OR EN.605.649

EN.705.743. ChatGPT from Scratch: Building and Training Large Language Models. 3 Credits.
Large language models (LLMs) like ChatGPT have ushered in a new wave of virtual assistants, chatbots, and text generators. Many see them as a paradigm shift in how humans interact with machines. Huge development ecosystems have arisen around LLMs, often abstracting away how they work to make them accessible to more people. While the democratization of this technology is important, LLMs cannot be fully harnessed and improved without understanding their inner workings at a fine level. In this course, students will build a small version of a text generation model like GPT3 over the course of several weeks. They will learn about the details of the GPT architecture from bottom to top, how the GPT architecture came about, and how it is used today in applications like ChatGPT. Once these fundamentals are established, students will build their own research experiment on top of their home-grown language models. Completing this course will prepare students to build and modify language models for further LLM research or novel applications.

EN.705.744. Computer Vision Using Transformers. 3 Credits.
Transformer networks are a new trend in Deep Learning. In the last decade, transformer models dominated the world of natural language processing (NLP) and have become the conventional model in almost all NLP tasks. However, developments of transformers in computer vision were still lagging. In recent years, application of transformers started to accelerate. This course will introduce the attention mechanism and the transformer networks by understanding the pros and cons of this architecture. The importance of unsupervised or semi-supervised pre-training for the transformer architectures, as well as their roles in foundation models will also be discussed. This will pave the way to introduce transformers in computer vision. Additionally, the course aims to will extend the attention idea into the 2D spatial domain for image datasets, investigate how convolution can be generalized using self-attention within the encoder-decoder meta architecture, analyze how this generic architecture is almost the same in image as in text and NLP, which makes transformers a generic function approximator, and discuss the channel and spatial attention, local vs. global attention among other topics. Further, time will be dedicated to studying the specific networks that are designed for mainstream computer vision tasks: classification, object detection and segmentation. In particular, ViT, shifted window transformer (Swin), Detection Transformer (DETR), segmentation transformer (SETR), and many others will be explored. The course concludes with the application of Transformers in video understanding with focus on action recognition and instance segmentation and will emphasize recent developments of transformers in large-scale pre-training and multimodal learning covering self-supervised learning, contrastive learning with masked image modeling, multimodal learning, and foundation CV models.
Prerequisite(s): EN.705.643 or equivalent PyTorch experience.

EN.705.801. Independent Study in Artificial Intelligence I. 3 Credits.
This course permits graduate students in artificial intelligence to work with a faculty mentor to explore a topic in depth or conduct research in selected areas. Requirements for completion include submission of a significant paper or project.Prerequisite(s): Seven artificial intelligence program graduate courses including the core courses, three elective courses. Students must also have permission of a faculty mentor, the student's academic advisor, and the program chair.

EN.705.802. Independent Study in Artificial Intelligence II. 3 Credits.
Students wishing to take a second independent study in artificial intelligence should sign up for this course.
Prerequisite(s): EN.705.801 Independent Study in Artificial Intelligence I and permission of a faculty mentor, the student's academic advisor, and the program chair.