

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.

Courses are offered online as well as in-person at the Applied Physics Laboratory.

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 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. To realize the impact of AI technologies requires a systems perspective that goes beyond the algorithms. 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, design, diagnostic, and deployment phases. The course will also cover ethics and bias in AI systems. The course includes a systems project that will encompass the full-lifecycle with interim milestones throughout the course. Homework assignments will be provided that in involves python programming.

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.705.621. 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 rapidly improving for well-defined tasks and narrow intelligence. Moreover, it has entered all industries in a myriad of ways. But will AI ever have human-like general intelligence? What does human-like general intelligence even mean? Why should we even care? This course is designed to answer these complex questions by giving students working knowledge of the underlying principles and mechanisms of human behavior and cognition, and how they may be applied to solving 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: ChatGPT 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