ARTIFICIAL INTELLIGENCE

The 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 comprehensive curriculum is designed to equip engineers and computer scientists with a solid grounding in the theoretical principles of artificial intelligence. Beyond theoretical understanding, students will gain practical knowledge and skills necessary to apply both existing and emerging AI concepts to real-world systems and processes. This dual focus ensures that graduates are well-prepared to implement AI solutions effectively in diverse settings and adapt to advancements in the field as AI technology evolves.

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

Erhan Guven, Assistant Program Manager
Senior Professional Staff
JHU Applied Physics Laboratory

Ian McCulloh, Assistant Program Manager
Senior Professional Staff
JHU Applied Physics Laboratory

Randal Burns, Bill and Lisa Stromberg Department Head
Professor of the JHU Computer Science Department
Johns Hopkins University

David A. Flanigan, Associate Dean for Non-Residential Graduate Education
Principal Professional Staff
JHU Applied Physics Laboratory

Conrad J. Grant, APL Chief Engineer
Principal Professional Staff
JHU Applied Physics Laboratory

Michael “Rabbi” Harasimowicz, Director
Artificial Intelligence Innovations
Lockheed Martin

John Hebeler, Instructor
Lockheed Martin Fellow
Johns Hopkins University

Anthony Hoogs, Vice President of AI
Kitware, Inc.

Kevin Ligozio, Chief AI Architect AOS
Principal Professional Staff
JHU Applied Physics Laboratory

Matteo G. Martemucci, Deputy Chief
Central Security Service
NSA

Gilbert Peterson, Professor of Computer Science
Air Force Institute of Technology

Jane Pinelis, Chief AI Engineer AIS Branch
Principal Professional Staff
JHU Applied Physics Laboratory

Christopher Ratto, AI Group Supervisor
Principal Professional Staff
JHU Applied Physics Laboratory

Benjamin Rodriguez, Program Co-chair Data Science
Principal Professional Staff
JHU Applied Physics Laboratory

Pedro Rodriguez, IS/CO Branch Supervisor
Principal Professional Staff
JHU Applied Physics Laboratory

Lawrence C. Schuette
Director of Research and Technology Programs
Lockheed Martin

David Silberberg, Program Chair ISE and Program Co-Chair RA&S
Principal Professional Staff
JHU Applied Physics Laboratory

Lanier Watkins, Program Chair CS, Program Chair CyS
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 OR EN.705.623

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. Production AI – Engineered AI Solutions. 3 Credits.
This course goes beyond theory, offering hands-on experience in building AI systems with the mindset and pace of a modern AI startup. Using a project-driven approach, students learn to architect, develop, and deploy real-world AI solutions entirely in the cloud, leveraging tools like Microsoft Azure, Terraform, and other cutting-edge technologies central to today's AI ecosystems. Students will incrementally build a production-grade, cloud-deployed AI system—individually and in teams—mirroring the end-to-end process of launching an AI startup. The course emphasizes not just tools, but the engineering mindset needed for building scalable, adaptable, and reliable AI systems. Key focus areas include: 1.) Data and Model Optimization – Streamlining data pipelines, adapting existing models, and using ensembles for efficiency and performance. 2.) System Integration – Developing distributed systems with messaging, NoSQL persistence, and robust monitoring. 3.) Cloud Deployment – Live updating through containerization and orchestration in a 100% cloud-based environment. By the end of the course, students will have built a portfolio-ready AI product and gained a deep, practical foundation in modern AI engineering for production.
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) provides a comprehensive overview of the foundational principles and techniques that empower machines to produce complex outputs, including text, images, video, and music. Students will examine the history and evolution of generative AI, tracing key milestones and landmark models that have shaped the field. The course begins with classical approaches —such as expert systems, genetic algorithms, Markov models, and constraint satisfaction problems—before advancing to modern generative techniques, including neural networks, autoencoders, generative adversarial networks (GANs), diffusion models, and Transformers. In addition to core models, students will be introduced to advanced topics such as prompt engineering, retrieval-augmented generation (RAG), and generative agents—autonomous AI systems capable of decision-making and goal-oriented behavior in dynamic environments. Ethical considerations and the societal impacts of generative AI, including concerns around bias, fairness, misinformation, and privacy, will be woven throughout the curriculum to foster responsible innovation. Assessment will include both research and project-based work, with students expected to design and implement a generative AI application from concept to deployment. By the end of the course, students will have a strong foundational understanding of generative AI and practical skills to creatively and ethically apply these technologies across diverse domains.

EN.705.606. Product Management for AI. 3 Credits.
This course provides a comprehensive overview of the principles and best practices for successful AI product management in real-world scenarios. Students will explore the end-to-end lifecycle of AI models, from problem identification through deployment, monitoring, maintenance, and continuous improvement. We will examine applications and constraints in government, healthcare, finance, and other industries. Topics include scoping, scalability, data availability, ethical considerations, and compliance requirements. We will also discuss bias, fairness, and unintended consequences of AI systems. Through case studies and hands-on team projects, students will learn how to identify potential AI solutions, manage risks, and ensure the positive impact of their AI products. Students will prepare product documentation suitable for use in the industry. By the end of the course, students will have developed a product pitch, a minimum viable product definition, and a risk assessment, preparing them to lead AI-driven innovation in their organizations.Prerequisites: Basic knowledge of machine learning, software engineering, or cloud computing is recommended.

EN.705.608. Applied Generative AI. 3 Credits.
This course provides a comprehensive exploration of generative artificial intelligence (AI) and its practical applications in solving complex business challenges. This course combines theoretical instruction with hands-on experience, equipping participants with the skills to design, develop, and deploy Generative AI-driven solutions. Topics include prompt engineering, natural language processing (NLP), generative AI workflows, fine-tuning large language models (LLMs), agentic AI development, secure AI practices, and the evaluation of AI solutions. Participants will also gain insight into ethical AI practices to ensure responsible and fair use of this technology. The curriculum includes two hands-on projects and analysis of real-world case studies.This program is designed for technology professionals, data analysts, engineers, consultants, technical managers, and STEM graduates seeking to enhance their expertise in generative AI. A foundational understanding of programming and data analysis is recommended. Participants will develop proficiency in using tools such as Python, VS Code, and various libraries. They will also learn techniques for retrieval-augmented generation (RAG), quick fine-tuning with vector databases, and use of open-source LLMs for proprietary applications. This course prepares graduates to apply generative AI to a wide range of applications, develop and train generative models, apply these techniques to create content, evaluate ethical implications, and analyze the impact of AI on various industries and society.

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.705.613. Responsible AI. 3 Credits.
This course explores the ethical, societal, and policy implications of artificial intelligence (AI) technologies, providing students with a comprehensive understanding of the responsibilities that come with their development and deployment. Through case studies, real-world examples, and critical discussions, students will examine issues related to fairness, transparency, privacy, accountability, and bias in AI systems. The course also delves into the regulatory landscape surrounding AI, exploring frameworks for responsible innovation and the governance structures needed to ensure AI aligns with human values. By the end of the course, students will be equipped with the tools to assess AI systems from an ethical standpoint, design responsible AI solutions, and contribute to the ongoing dialogue around the future of AI and its impact on society.

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.705.617. Artificial Intelligence in Healthcare. 3 Credits.
This course equips students with the applied knowledge and engineering mindset needed to design and deploy AI solutions in complex healthcare environments. Structured around real-world healthcare workflows, students will explore how artificial intelligence can be integrated to enhance clinical decision-making, improve operational efficiency, and support patient outcomes. Through hands-on projects, learners engage with tools and techniques used in modern healthcare AI systems, from predictive modeling and clinical decision support to robotic process automation and the responsible use of large language models. Students begin by developing a foundational understanding of AI technologies within the healthcare lifecycle, including regulatory and ethical frameworks. They then build and evaluate AI models that augment clinical reasoning, such as risk scoring and diagnostic support. The course also explores the application of generative AI to medical documentation and patient communication and culminates in the implementation of robotic and process automation tools to streamline healthcare workflows. Throughout the course, students will iteratively develop a portfolio-ready healthcare AI solution, working both individually and in teams. By the end, they will have a deep, practical understanding of how to translate AI capabilities into impactful, ethical, and integrated applications within the healthcare system.

EN.705.618. Neuromarketing AI. 3 Credits.
In today's rapidly evolving marketplace, the key to effective leadership is understanding the intersection of science, technology, and consumer behavior. Neuromarketing AI: The Future of Persuasion and Consumer Insights, is a pioneering course designed specifically for leaders who seek to harness the latest advancements in neuroscience and artificial intelligence to transform their marketing strategies and decision-making processes. This course merges groundbreaking brain science with AI-driven techniques, giving participants a unique opportunity to master the science of persuasion and its real-world applications. Participants will explore how the human brain processes decisions, learn how to apply neuromarketing principles, and leverage AI to predict consumer behavior and personalize messaging at scale. With a curriculum built on the renowned Persuasion Code model, this program will empower participants to revolutionize how they communicate with customers, employees, and stakeholders.

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.705.623. AI Algorithm Design and Analysis. 3 Credits.
This course provides an in-depth exploration of the design, analysis, and optimization of algorithms critical to artificial intelligence. Building on foundational algorithmic techniques such as dynamic programming, graph theory, and heuristic search, the course extends into applications across machine learning, natural language processing, computer vision, and generative AI. Topics include computational complexity, probabilistic reasoning, optimization methods, and trade-offs in algorithm performance. Students will gain theoretical and practical insights essential for solving real-world AI problems, including designing hybrid AI systems that integrate reasoning, optimization, and pattern recognition.
Prerequisite(s): EN.605.202 Data Structures

EN.705.625. Introduction to Agentic AI. 3 Credits.
In this course, we will dive into what it truly means to build agentic AI —systems that perceive their environment, make decisions, learn from experience, and act autonomously. We begin by exploring the core principles of intelligent agency and how these systems interact with the world around them. From there, we will examine foundational models and techniques for crafting intelligent behavior, including decision trees, utility theory, Markov decision processes, and game theory. As the course progresses, you will design both standalone agents and complex multi-agent systems, learning how they make decisions, communicate, collaborate, and compete. We will also explore the human side of AI—how to build trust, ensure explainability, and design intuitive interfaces for effective human-AI collaboration. In the final phase, we investigate the frontier of AI: generative agents powered by large language and vision models. These advanced systems do not merely react—they reflect, plan, and interact in ways that resemble human cognition. You will gain hands-on experience building and evaluating these agents using state-of-the-art frameworks. By the end of the course, you will be equipped to model, build, and evaluate intelligent agents for both simulated and real-world environments, with a solid grounding in both theory and practice.

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.705.641. Natural Language Processing: Self-Supervised Models. 3 Credits.
In recent times, Large Language Models (LLMs) have earned the attention of the world. OpenAI's infamous generative LLM, ChatGPT, became the fastest-growing consumer application in history in only two months– and the feverish interest around LLMs continues to grow. Large self-supervised (pre-trained) models (such as LLMs) 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 analyze their own self-supervised neural network models using the PyTorch framework.

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. Deep Learning 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, applications 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 impact for developments of large-scale foundation model. This will pave the way to introduce transformers in computer vision. Additionally, the course aims to 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. Furthermore, we will also study different neural architectures that are designed for several fundamental tasks in computer vision, namely, classification, object detection, semantic and instance segmentation. In particular, vision transformer, pyramid vision transformer, shifted window transformer (Swin), Detection Transformer (DETR), segmentation transformer (SETR), and many others will be explored. The course also examines 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 vision foundation 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.