## **ARTIFICIAL INTELLIGENCE ENGINEERING (AIE)**

AIE 355: Optimization for Machine Learning

## 3 Credits

This course provides an overview of modern optimization methods, tailored specifically for applications in the rapidly evolving fields of machine learning and data science. As the complexity of models and the scale of datasets continue to grow exponentially, the ability to leverage efficient and scalable optimization techniques has become an indispensable skill. The curriculum covers a wide range of advanced optimization algorithms, including (stochastic) gradient-based methods, momentum techniques, adaptive methods, constrained optimization, and specialized techniques for large-scale and distributed computing environments. Particular emphasis is placed on the scalability of these algorithms, ensuring that students gain a thorough understanding of how to effectively handle massive datasets and high-dimensional problems. Through a combination of theoretical foundations and handson implementations in computer-based homework exercises, students will learn to navigate the intricate landscape of optimization challenges, from training deep neural networks and tuning hyperparameters to solving complex constrained optimization problems. By the end of this course, students will acquire a comprehensive toolkit of optimization methods, along with the ability to critically evaluate and select the most appropriate techniques for their specific machine learning and data science applications.

Enforced Prerequisite at Enrollment: Prerequisites: MATH 225 or (MATH 220 and (MATH 231 or MATH 230) and (CMPSC 131 or (CMPSC 121 and DS 120)))