Learning Objectives: Python for Cloud & AI ("Python 200")

Last updated Spring 2026

Note: *Python for Cloud & AI is still in development and will run a pilot cohort starting in March 2026. The course learning objectives and cadence may change slightly after the pilot.*

Course Objectives

By the end of this course, students will:

1. **Apply statistical thinking and build production-ready data pipelines**
    - Analyze datasets using statistical measures and hypothesis testing
    - Design and implement automated data analysis workflows using pipeline orchestration tools
2. **Build and deploy machine learning solutions**
    - Train, evaluate, and fine-tune machine learning models for regression, classification, and computer vision tasks
    - Apply transfer learning techniques to adapt pretrained models for custom applications
    - Evaluate model performance and identify issues like overfitting and bias
3. **Develop AI-powered applications**
    - Integrate large language models through commercial APIs to build chatbots and interactive applications
    - Implement retrieval-augmented generation (RAG) systems using vector databases to supplement LLM knowledge
    - Design and build AI agents that autonomously plan, reason, and execute multi-step tasks
4. **Deploy scalable cloud-based systems**
    - Provision and configure cloud infrastructure (compute, storage, databases) using Azure
    - Deploy machine learning models and AI applications on cloud platforms
    - Build end-to-end ETL pipelines using both low-code and code-first approaches

Course Structure

- Statistical thinking (1 week)

- ML (3 weeks)
- AI (3 weeks)
- Cloud (3 weeks)