Machine Learning and AI Essentials with Python

- 3 Days
- Lecture and Hands-on Labs

Course Overview

Machine learning has evolved rapidly with the rise of large language models, AI-assisted development, and modern data pipelines. This three-day hands-on course teaches students how to build practical machine learning solutions using Python while integrating modern AI tools and workflows. Participants begin with a Python and data science refresher before exploring core machine learning techniques including regression, classification, and clustering. The course then moves into building complete machine learning pipelines, evaluating model performance, and improving models using ensemble techniques. Modern AI topics are woven throughout the course including AI-assisted coding, prompt-driven data analysis, and integrating large language models into Python applications. Students complete several hands-on labs where they build models, analyze datasets, visualize results, and integrate AI services into applications. By the end of the course, students will understand how to design, build, evaluate, and deploy practical machine learning solutions using Python and modern AI tooling.

Review this course online at https://www.alta3.com/courses/ai-ml

Who Should Attend

- Python Developers expanding into Machine Learning
- Data Analysts working with predictive models
- Engineers working with automation and AI systems
- Data Scientists seeking stronger ML pipeline skills

What You'll Learn

- Design complete machine learning pipelines.
- Implement core ML techniques such as regression, classification, and clustering.
- Integrate AI services and LLMs into applications.
- Evaluate and enhance model performance with ensemble methods.

Outline

Python Foundations for Machine Learning

1. Review of Python fundamentals for data analysis
2. Working with lists, dictionaries, and arrays
3. Python workflows for data science projects #### Working with Jupyter Notebooks and ML Environments
4. Setting up Anaconda and Jupyter
5. Notebook workflow best practices
6. Running Python experiments interactively #### Core Python Libraries for Machine Learning
7. NumPy for numerical computing

8. pandas for dataset manipulation
9. matplotlib and seaborn for visualization #### Introduction to Artificial Intelligence and Machine Learning
10. Understanding the AI landscape
11. Differences between ML, deep learning, and LLMs
12. Common AI use cases across industries #### Machine Learning Concepts and Workflows
13. Training data and features
14. Labels and target variables
15. Model training and prediction #### Data Preparation and Feature Engineering
16. Data cleaning techniques
17. Handling missing data
18. Managing inconsistent datasets #### Feature Engineering Concepts
19. Feature scaling
20. Encoding categorical variables
21. Detecting and handling outliers #### Supervised Machine Learning
22. Linear regression fundamentals
23. Multiple regression models
24. Practical regression use cases #### Classification Models
25. Binary vs multi-class classification
26. Logistic regression
27. Decision trees #### Unsupervised Machine Learning
28. Understanding clustering
29. K-means clustering
30. Real world clustering applications #### Data Visualization for Machine Learning
31. Scatter plots and distribution charts
32. Correlation analysis
33. Feature importance visualization #### Model Evaluation and Optimization
34. Training vs testing data
35. Accuracy, precision, recall, and F1 score
36. Avoiding overfitting and underfitting #### Ensemble Learning
37. Random forests
38. Gradient boosting concepts
39. Ensemble model advantages #### Machine Learning Pipelines
40. Data ingestion
41. Feature engineering
42. Model training and evaluation #### Integrating Generative AI into Python Applications
43. Overview of LLM architectures
44. Using APIs for AI services
45. AI-assisted development workflows #### Explainable AI and Responsible AI
46. Feature importance
47. Model transparency #### Final Hands-On Project
48. Prepare a dataset
49. Train a machine learning model
50. Evaluate performance
51. Integrate AI features into the workflow

Labs

- Exploring and Analyzing a Dataset with Python
- Exploring Machine Learning Libraries
- Preparing a Dataset for Machine Learning
- Building a Regression Model in Python
- Implementing Classification Models
- Performing Customer Segmentation with Clustering

- Visualizing Model Results
- Evaluating Model Performance
- Building Ensemble Models
- Creating a Complete Machine Learning Pipeline
- Using Python to Interact with an LLM API
- Visualizing Model Explainability
- End-to-End Machine Learning Project

Prerequisites

- Basic Python programming knowledge
- Familiarity with Python data structures
- Basic understanding of statistics and problem solving
- Prior exposure to pandas or numpy is helpful but not required

Certification

- Python Artificial Intelligence (AI) for the Enterprise - Certification Project