Machine Learning

Level: Introductory

Introduction to Machine Learning: Foundations to Implementation

3 days

Welcome to this comprehensive introduction to machine learning. This course provides a balanced approach to understanding and implementing machine learning solutions, combining essential theoretical knowledge with practical, hands-on experience.

Machine learning is transforming industries across the global economy, with ML systems driving nearly $4 trillion in business value annually. As companies increasingly adopt AI-powered applications, professionals with the ability to design, implement, and evaluate machine learning systems are in unprecedented demand.

This course takes you through the complete machine learning lifecycle, from data preparation to model deployment. You'll gain proficiency in both supervised and unsupervised learning techniques, understand the mathematical principles that underpin effective models, and learn industry best practices for building production-grade ML systems.

What sets this course apart is our focus on real-world implementation. Each module balances conceptual understanding with practical application through carefully designed labs and case studies drawn from actual business scenarios. By the end of the three days, you'll not only understand how machine learning works but be able to apply these techniques to your own data challenges.

Learning Outcomes

By the end of this course, participants will be able to:

- Understand the fundamental concepts and mathematics behind machine learning algorithms
- Prepare and preprocess datasets for effective model training
- Implement and evaluate supervised learning models for regression and classification tasks
- Apply unsupervised learning techniques for clustering and dimensionality reduction
- Optimize models through systematic evaluation and hyperparameter tuning
- Understand the basics of neural networks and deep learning applications
- Deploy machine learning models in production environments
- Monitor and maintain ML systems throughout their lifecycle

Course Outline

Module 1: Machine Learning Landscape

- Contrast between traditional programming and machine learning approaches
- Overview of key learning paradigms: supervised, unsupervised, and reinforcement learning
- The machine learning lifecycle from business goal identification to model monitoring
- Introduction to popular ML frameworks and libraries in the Python ecosystem
- Current applications and business impact of machine learning across industries

Module 2: Mathematical Foundations

- Essential linear algebra concepts for machine learning (vectors, matrices, operations)
- Calculus fundamentals for optimisation algorithms
- Probability and statistics principles underlying ML algorithms
- The role of mathematical intuition in selecting appropriate algorithms
- Case study: How mathematical misconceptions lead to model failures in production

Module 3: Data Preprocessing Techniques

- Strategies for handling missing values, outliers, and data quality issues
- Feature engineering approaches to improve model performance
- Normalisation and standardisation techniques for different algorithm types
- Categorical encoding methods and their impact on model performance
- Addressing class imbalance issues in classification problems

Module 4: Regression Analysis

- Linear regression fundamentals and assumptions
- Regularisation techniques (L1/L2) to prevent overfitting
- Tree-based regression methods for handling non-linear relationships
- Evaluation metrics for regression models and their business interpretations
- Lab: Building and evaluating a production-ready regression pipeline

Module 5: Classification Systems

- Logistic regression for binary and multi-class problems
- Decision trees and ensemble methods (Random Forests, Gradient Boosting)
- Support Vector Machines and kernel selection strategies
- Imbalanced classification challenges in fraud detection and medical diagnosis

- Case study: Precision-recall trade-offs in critical decision systems

Module 6: Model Evaluation and Optimisation

- Cross-validation strategies for robust performance estimation
- Hyperparameter tuning approaches and automation techniques
- Learning curves for diagnosing bias-variance issues
- Feature selection methods to improve model efficiency
- Lab: Systematic optimisation of an ML model with measurable performance gains

Module 7: Unsupervised Learning Applications

- Clustering algorithms (K-means, DBSCAN) and their applications
- Dimensionality reduction techniques (PCA, t-SNE) for visualisation and preprocessing
- Association rule mining for market basket analysis
- Anomaly detection approaches for identifying outliers and unusual patterns
- Lab: Implementing customer segmentation using unsupervised learning

Module 8: Deep Learning Fundamentals

- Neural network architecture and components
- Activation functions and their impact on model performance
- Backpropagation and gradient-based optimisation
- Transfer learning for leveraging pre-trained models
- Practical considerations when implementing neural networks

Module 9: Deployment Lifecycle Management

- Model packaging and containerisation for deployment
- Implementing CI/CD pipelines for ML systems
- Monitoring frameworks for tracking model performance in production
- Strategies for handling concept drift and model degradation
- Lab: Deploying a complete machine learning API with monitoring

Conclusion and Next Steps

Throughout this course, you'll build a solid foundation in machine learning that balances theoretical understanding with practical implementation skills. The concepts and techniques covered will equip you to tackle real-world data problems and integrate machine learning capabilities into your organisation's technology stack.

By completing this course, you'll join a growing community of professionals who can leverage the power of machine learning to drive business value. You'll leave with not only new technical

skills but also the ability to identify which problems are amenable to machine learning solutions and how to approach them systematically.

As machine learning continues to transform industries, the knowledge gained in this course will serve as a valuable foundation for your ongoing professional development in this dynamic field.

Intended Audience

This course is designed for software developers, data analysts, and IT professionals who want to incorporate machine learning into their skill set. It's ideal for those with programming experience who need to understand both the theoretical foundations and practical applications of machine learning in business contexts.

Prerequisites

Those attending this course should meet the following:

- Basic Python programming knowledge
- Familiarity with fundamental data structures and algorithms
- Understanding of basic statistics concepts (mean, median, standard deviation)
- Experience working with data in tabular formats
- Comfort with command-line operations for environment setup