Machine Learning

Level: Intermediate

AI/ML for Data Scientists: Modern Approaches and Deployment

3 days

AI/ML for Data Scientists: Modern Approaches and Deployment

You've mastered the basics of machine learning - now it's time to dive into the cutting edge. This course is your gateway to the most exciting developments in AI, from foundation models that can understand natural language to sophisticated neural architectures that can process multiple types of data simultaneously. We'll take you beyond theoretical concepts to show you how to implement these powerful technologies in production environments.

The world of AI is moving at an incredible pace, and as a data scientist, you're at the forefront of this revolution. This course is designed to give you both the theoretical understanding and practical skills needed to implement modern AI approaches that solve real business problems. We'll focus on the techniques and architectures that are transforming the field, while ensuring you can deploy them reliably and responsibly.

What distinguishes this course is its focus on bridging theory with practice. Each module combines essential concepts with hands-on implementation, ensuring you not only understand how these technologies work but can apply them to solve real-world challenges. We'll move beyond the basics of neural networks to explore advanced architectures, self-supervised learning, and state-of-the-art approaches for different data types.

Beyond model building, we'll address the critical aspects of bringing AI systems to production through robust MLOps practices. You'll learn how to deploy, monitor, and maintain AI models while adhering to ethical principles and governance frameworks that ensure responsible AI implementation.

Learning Outcomes

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

- Implement advanced deep learning models for diverse data types
- Fine-tune and deploy large language models for domain-specific applications
- Build and evaluate multimodal AI systems that combine textual and visual information
- Design robust MLOps pipelines for enterprise AI solutions
- Monitor and maintain ML systems in production environments
- Apply transfer learning and foundation models to practical business problems
- Navigate the ethical considerations of advanced AI deployments

- Implement retrieval-augmented generation systems for knowledge-intensive applications

Course Outline

Module 1: AI/ML Landscape and Deep Learning Fundamentals

- Evolution of AI/ML approaches and their business applications
- Neural network fundamentals: architectures, components, and training mechanics
- Mathematical foundations essential for deep learning implementation
- Transitioning from traditional ML to neural network approaches
- Framework selection and setup: PyTorch vs. TensorFlow ecosystems
- Lab: Implementing and optimising basic neural networks for tabular data

Module 2: Transfer Learning and Foundation Models

- Understanding the foundation model revolution and its implications for data science
- Transfer learning principles and implementation techniques
- Strategies for effective fine-tuning and domain adaptation
- Embedding extraction and utilisation from pre-trained models
- Evaluating transfer learning effectiveness across different domains
- Lab: Fine-tuning pre-trained models for custom classification tasks

Module 3: Advanced Neural Network Architectures

- Convolutional neural networks for image and spatial data processing
- Recurrent networks and their applications for sequential data
- Attention mechanisms and their role in modern architectures
- Modern architectures for structured and unstructured data
- Best practices for architecture selection based on problem type
- Lab: Implementing CNN-based systems for domain-specific applications

Module 4: Self-Supervised Learning Approaches

- Paradigm shift from supervised to self-supervised learning
- Contrastive learning techniques and implementation strategies
- Representation learning for unlabelled data utilisation
- Domain-specific self-supervised pre-training approaches
- Evaluating the quality of learned representations
- Lab: Building a self-supervised learning pipeline for business data

Module 5: Large Language Models (LLMs)

- Architecture and capabilities of transformer-based language models
- Understanding attention mechanisms and their optimisation
- Prompt engineering techniques for optimal LLM utilisation
- Parameter-efficient fine-tuning approaches (LoRA, P-Tuning, QLoRA)
- Optimising inference for production deployment
- Lab: Fine-tuning smaller LLMs for domain-specific applications

Module 6: Multimodal AI Architectures

- Fundamentals of multimodal learning and business applications
- Recent advances in multimodal large language models
- Architectural approaches for combining different data modalities
- Projection and alignment techniques for multimodal systems
- Evaluation frameworks for multimodal model performance
- Lab: Implementing a multimodal system for business document analysis

Module 7: Retrieval-Augmented Generation (RAG)

- RAG architecture and its applications in enterprise contexts
- Vector database selection and optimisation
- Chunking strategies for effective knowledge retrieval
- Hybrid search approaches and re-ranking techniques
- Evaluation metrics for RAG system quality assessment
- Lab: Building and optimising a RAG system for company knowledge bases

Module 8: MLOps Fundamentals

- MLOps lifecycle and key components for enterprise implementations
- Version control strategies for data, models, and experiments
- Continuous integration and deployment for ML systems
- Infrastructure considerations for scalable AI deployments
- Governance frameworks for enterprise AI systems
- Lab: Implementing a complete ML pipeline with appropriate versioning

Module 9: Model Deployment and Monitoring

- Deployment architecture selection for different use cases
- Containerisation and orchestration for AI systems

- Monitoring strategies for production ML models
- Drift detection and automated retraining frameworks
- A/B testing methodologies for model updates
- Lab: Deploying models with monitoring and automated alerting

Module 10: Responsible AI and Ethical Frameworks

- Ethical considerations in advanced AI implementations
- Bias detection and mitigation techniques
- Model interpretability and explainability methods
- Regulatory compliance and governance frameworks
- Privacy-preserving machine learning approaches
- Lab: Conducting comprehensive ethical assessments of AI systems

Conclusion and Next Steps

This course provides a comprehensive foundation in modern AI/ML approaches specifically tailored for data scientists looking to elevate their practice. The combination of theoretical knowledge and practical implementation skills will enable you to design, develop, and deploy sophisticated AI systems that address real business challenges.

As the field continues to evolve rapidly, we'll discuss strategies for staying current with emerging techniques and tools. You'll leave this course with not only technical skills but also a framework for evaluating new AI approaches and determining their applicability to your specific domain.

Intended Audience

This course is designed for practicing data scientists, machine learning engineers, and technical professionals who want to upgrade their skills with modern AI approaches. It's ideal for those who have experience with traditional machine learning methods and want to master deep learning, large language models, multimodal AI, and MLOps practices for enterprise-grade deployments.

Prerequisites

Those attending this course should meet the following:

- Intermediate Python programming skills
- Basic understanding of machine learning concepts and workflows
- Familiarity with data manipulation libraries (Pandas, NumPy)
- Working knowledge of statistical concepts
- Experience with building and evaluating basic machine learning models