DATA SCIENCE | ADVANCED | 5 DAYS

AI and MLOps Fundamentals

Build a strong foundation in designing, developing, and deploying AI and ML models from initial scoping to production. Learn how to optimize model performance, address scaling challenges, and collaborate across teams to deliver reliable, high ROI AI solutions.

Why AI/MLOps?

Enhanced Operational Efficiency

AI and ML will generate US$4.4 trillion in business value by next year! Upskilling your teams in these areas will help in boosting their productivity and streamline your workflows, significantly boosting overall operational efficiency.

Reduced Operational Risks

This AI and ML Ops training covers risk mitigation modules, helping your employees identify and bridge the gaps in building and deploying AI/ML models at scale and within modest durations.

Improved Decision-Making

Training in AI and ML Ops sharpens employees' skills in working with and extracting valuable insights from complex data sets. This not only prepares them to make informed, precise decisions but also empowers them to leverage advanced analytical tools for better business outcomes.

Future-Agnostic Innovation

A workforce trained in next-gen technologies like AI and ML can innovate new solutions that make your business adaptable to tech advancements, staying competitive in an evolving tech ecosystem.

What you'll be able to do

- ✓ Learn to design and implement AI/ML models from initial scoping through to deployment
- ✓ Identify and rectify gaps in developing and scaling these models
- ✓ Master techniques to evaluate and optimize AI/ML models for specific projects
- ✓ Improve interdisciplinary team collaboration within AI projects
- ✓ Learn how to strategize for reliability, scalability, and performance, maximizing ROI in AI/ML initiatives

Who this is for

- AI and Data Science practitioners
- IT professionals planning to transition to AI/ML development
- Operations professionals planning to upgrade to AI and MLOps

Curriculum

01 Maths and Programming Pre-requisites

- Fundamentals of statistical analysis and probability theory
- Calculus and linear algebra with an emphasis on tensors
- Basics of Python with a focus on data science
- TensorFlow for implementing tensor operations
- Data munging techniques for preprocessing and transforming tabular data

02 Foundations of Machine Learning and Artificial Intelligence

- How to approach problems using data and algorithms in AI/ML workflows
- Data types and models:
- Gradient-boosted models for structured datasets
- Convolutional Neural Networks for image processing
- Recurrent Neural Networks for sequential data analysis
- Transformers and GPT models for text processing
- Reinforcement Learning for robotics
- Deep Learning fundamentals
- Shallow neural networks - Sigmoidal NNs and Logistic Regression
- Network training techniques including Momentum, AdaGrad, and ADAM
- Practical implementation with Keras

03 Computer Vision

- Fundamental use cases
- Convolutional operations including the use of kernels, padding, and feature maps
- Pooling techniques for reducing spatial dimensions
- CNNs for image classification
- Transfer learning to leverage pre-trained models for new tasks
- Residual connections and batch normalization for deeper neural networks
- Convolution methods - depthwise separable convolution and the Xception architecture
- Object localization and detection using YOLO
- Image segmentation with UNet and DeepLab
- Best practices and tools

04 Designing Machine Learning Systems

- Basics of cloud computing
- Introduction to AWS and Azure, with a focus on specific services - Cloud9 and EC2 instances
- AWS SageMaker for deploying ML models
- Cloud computing tools - AWS ML Tools, Google Cloud, and IBM Watson

05 Practical MLOps

- Fundamentals of DevOps, MLOps, Continuous Integration (CI) and Continuous Deployment (CD)
- Basics of version control with Git and common commands like clone, pull, push, branch, commit & resolving merge conflicts
- How to implement containerization with Docker and Docker Compose files for managing multi-container setups
- CI tools and Data Version Control (DVC) with GitHub Actions
- MLOps tools including Jenkins, DagsHub, Weights, and Biases
- Basics of Python Flask

06 Natural Language Processing (NLP)

- Foundations of NLP
- Data preprocessing techniques for language models - Text Vectorization Layer, Standardization, Vocabulary Indexing, Embedding Word Vectors, TF-IDF
- Basics of Tokenization
- Bag of Words (BoW)
- Attention mechanism and transformer encoders for comprehension
- Architecture and applications of BERT

07 Representation Learning, Generative Models, and Research Trends

- Fundamentals of Representation Learning
- Decoder-only GPT class models
- LLMOps vis-a-vis tools and platforms such as LangChain and the OpenAI API
- Basics of prompt engineering, RAG (Retriever-Augmented Generation), and LoRA (Low-Rank Adaptation)
- Generative Adversarial Networks (GANs) and other diffusion models
- Autoencoders and their use in pre-training models
- How to implement Reinforcement Learning through Human Feedback (RLHF)
- Guardrails and processing techniques for LLM outputs

08 Parallel Computer Architecture and Programming Models

- Fundamentals of Parallel Programming Paradigms via OpenMP
- Applications of OpenMP for parallel algorithms - k-means clustering, histogram computation
- Other parallel programming paradigms such as MPI and CUDA
- Submitting jobs on computing clusters
- Comparing Python parallel programming - Multi-Threaded vs Multi-Processor
- Introduction to PyOMP and PythonMPI

09 Machine Learning at Scale

- Utility of GPUs in ML through architecture essentials and the specific roles of tensor cores
- Distributed computing techniques using TensorFlow
- Introduction to Parallel ML libraries including Dask, Horovod, Ray, & RAPIDS, and GPU-accelerated libraries such as cuDF, cuML, and cuPY