

About the Company

We aim to bridge the gap between academic learning and industry expectations through practical, industry-oriented education. We are building a community of industry professionals passionate about sharing their expertise and making high-quality, industry-relevant education accessible to aspiring professionals.

Benefits of Completing the Course

- **Production-Ready AI Portfolio:** Graduate with enterprise-grade capstone projects spanning core ML, MLOps, and Generative AI on GitHub.
- **Modern Tech Stack Mastery:** Gain hands-on expertise in Python, PyTorch, Scikit-Learn, FastAPI, Docker, DVC, MLflow, Airflow, Vector DBs, and LLM orchestration.
- **Real-World Engineering:** Solve actual business challenges in predictive modeling, computer vision, RAG systems, and production model monitoring.
- **Career Acceleration:** Position yourself strongly for high-demand roles like AI/ML Engineer, MLOps Engineer, and GenAI Engineer.

Month-by-Month Course Overview

Month 1: Python, Data Engineering & EDA Foundations

- **Module 1: Python for AI/ML & Data Handling:** NumPy vectorized computing and array broadcasting; Pandas DataFrames, indexing, filtering, and aggregations; Real-world data ingestion (CSV, JSON, SQL, Parquet); Memory optimization and reusable preprocessing script design.
- **Module 2: Exploratory Data Analysis & Feature Engineering:** Automated data profiling and distribution analysis; Categorical encoding (One-Hot, Label, Target, Ordinal); Feature scaling (MinMax, Standard, Robust); Outlier management, detection, and capping strategies.
- **Module 3: Advanced Feature Engineering & Unsupervised Learning:** Advanced data imputation and automated pipeline preprocessing; Dimensionality reduction using PCA and t-SNE; Unsupervised clustering (K-Means, DBSCAN, hierarchical); Anomaly detection and handling imbalanced datasets (SMOTE).

Month 2: Applied Machine Learning & Evaluation

- **Module 1: Supervised Learning & Scikit-Learn Workflows:** Modular Scikit-Learn ML pipelines; Linear and Logistic Regression implementation; Tree-based models (Decision Trees, Random Forests, Extra Trees).
- **Module 2: Gradient Boosting Frameworks:** Extreme Gradient Boosting with XGBoost; LightGBM optimization for large tabular data; CatBoost for categorical feature handling.
- **Module 3: Model Evaluation, Validation & Tuning:** Classification metrics (Precision, Recall, F1-Score, ROC-AUC); Regression metrics (RMSE, MAE, R-Squared); Stratified K-Fold and Time-Series cross-validation splits; Hyperparameter optimization with Optuna and experiment tracking via MLflow.

Month 3: Deep Learning Foundations & Mid-Course Capstone

- **Module 1: Deep Learning Foundations with PyTorch:** PyTorch tensors, CUDA acceleration, and Autograd engine; Building Multi-Layer Perceptrons (MLPs) from scratch; Loss functions, optimizers (Adam, AdamW, SGD), and training loops; Model checkpointing, early stopping, and weight persistence.
- **Module 2: Computer Vision Systems:** Convolutional Neural Networks (CNNs) architecture; Image processing with OpenCV and Albumentations; Transfer learning with pre-trained models (ResNet, EfficientNet); Real-time object detection using YOLO and bounding box rendering.
- **Module 3: NLP, Sequential Models & Mid-Course Capstone:** Text preprocessing, tokenization, and vector embeddings (Word2Vec, TF-IDF); Sequential models (RNNs, LSTMs) and modern transformers (BERT, RoBERTa); Hugging Face transformers ecosystem pipelines; *Mid-Course Capstone Project:* End-to-end core ML, CV, or NLP application build and evaluation.

Month 4: Production MLOps, Model Serving & Data Versioning

- **Module 1: Model Packaging & APIs:** Model serialization (Pickle, Joblib, ONNX, TorchScript); Asynchronous REST inference APIs built with FastAPI; Pydantic schemas, validation, and error handling; Inference latency reduction and memory footprint optimization.
- **Module 2: Containerization & Data Version Control:** Docker containerization for machine learning microservices; Data Version Control (DVC) for datasets and model artifacts; Reproducible pipeline DAG creation using dvc.yaml.
- **Module 3: Cloud Deployment & Optimization:** Cloud model serving on AWS SageMaker and GCP Vertex AI; Model compression via quantization and weight pruning; Accelerated inference execution using ONNX Runtime; Production security, secrets management, and environment isolation.

Month 5: Cloud Deployment, Scalable Pipelines & Monitoring

- **Module 1: Workflow Orchestration & Feature Stores:** Offline and online feature management using Feast; Pipeline scheduling and orchestration via Prefect / Airflow; Automated schema validation checks before pipeline execution.
- **Module 2: MLOps & Production Monitoring:** Identifying Data Drift and Concept Drift in production environments; Continuous monitoring and reporting with Evidently AI; Centralized logging, metric collection, and failure alerting strategies; Automated trigger-based vs. scheduled model retraining policies.

Month 6: Industry-Grade GenAI, Vector Databases & Final Capstone

- **Module 1: Generative AI Foundations & RAG Architectures:** Foundation model landscape (OpenAI, Anthropic, Ollama); Advanced prompting (Chain-of-Thought, ReAct framework, Pydantic guardrails); Vector search distance metrics (Cosine, L2) and semantic chunking; LangChain, LlamaIndex orchestration, and hybrid search (BM25 + vector).
- **Module 2 & 3: Vector DBs, Fine-Tuning, Streaming & Final Capstone:** Enterprise vector stores (Pinecone, Milvus, Qdrant) and ANN indexing (HNSW); PEFT, LoRA, QLoRA via Hugging Face TRL; vLLM serving; Multi-agent frameworks (CrewAI, LangGraph) & Ragas evaluation; Real-time streaming with Kafka/Redpanda, GitHub Actions CI/CD, OWASP Top 10 for LLMs; *Final Industry Capstone Project:* End-to-end enterprise-grade GenAI or MLOps production deployment.

Learning Journey Visual Roadmap (Step-by-Step Process Flow Layout)

1

Month 1 Roadmap: Python, Data Engineering & EDA
Deliverable: NumPy/Pandas data pipelines & automated EDA profiles

2

Month 2 Roadmap: Applied Machine Learning & Tuning
Deliverable: Scikit-Learn models, XGBoost, and MLflow experiment tracking

3

Month 3 Roadmap: Deep Learning & Mid-Capstone
Deliverable: PyTorch neural networks, YOLO computer vision, & Mid-Course Capstone

4

Month 4 Roadmap: MLOps, Serving & Versioning
Deliverable: FastAPI microservices, Docker containers, & DVC data versioning

5

Month 5 Roadmap: Orchestration & Drift Monitoring
Deliverable: Airflow/Prefect pipelines, Feast feature store, & Evidently AI drift alerts

6

Month 6 Roadmap: GenAI, Vector DBs & Final Capstone
Deliverable: Enterprise RAG systems, fine-tuned LLMs, vLLM serving, & Final Industry Capstone

Course Pricing & Investment

Course Name	Original Price	Discount	Price After Discount
First 3 Months (Core ML)	₹50,000	30%	₹35,000
Next 3 Months (Advanced ML)	₹50,000	30%	₹35,000
6-Month Course (Full Course)	₹70,000	30%	₹49,000