

AI-NATIVE PROFESSIONAL CERTIFICATE PROGRAMME

Data Science & AI Programme

A nineteen-week path from statistics and Python through classical machine learning, deep learning, generative AI and MLOps – ending with a deployed, monitored model, not just a trained one in a notebook.

Beyond Notebooks

Every learner ships a model as a live, monitored API – not just a trained .ipynb file.

GenAI & LLMs Built In

RAG, LangChain and vector databases are a core phase, not a bolt-on elective.

Live, Mentor-Led

Taught personally by a practising Senior Data Scientist – not pre-recorded video.

Eight Capstone Tracks

Choose a capstone in your own domain, from dynamic pricing to a RAG chatbot.

19

Weeks

95

Sessions

1.5 hrs

Per session

Mon-Fri

Live schedule

34%

Projected growth in data scientist roles through 2034, according to the US Bureau of Labor Statistics – among the fastest-growing occupations tracked. **This programme is built for where that growth is headed.**

Learn directly from someone who has **done** the work



Aravind Reddy Dudipala

Founder, Unicus Learning
Senior Data Scientist & AI/ML
Programme Leader

11 Years Experience

Berkeley Haas ML/AI Instructor

**Imperial College Business School
Instructor**

Aravind founded Unicus Learning and personally leads its training programmes in data analytics, data science, Power BI and agentic AI for working professionals. He brings **11 years of experience** as a Senior Data Scientist and AI/ML Programme Leader, and works with international universities, teaching Professional Certification courses at **Berkeley Haas** and the **Imperial College Business School**. Alongside this, he mentors working professionals across the globe in AI, machine learning, data science and data analytics.

"A model that only exists in a notebook is a prototype. A model behind a monitored API is a product. I teach people to build the second one."

That philosophy is why this programme's deployment and generative AI phases are treated as core curriculum, not a bonus module at the end – and why the capstone requires a working, deployed system rather than a final notebook.

Every batch – up to 50 learners – is taught personally and live, session by session. No handoff to a teaching assistant.

WHAT WORKING WITH ARAVIND LOOKS LIKE

- Founder of Unicus Learning, personally involved in every batch he teaches
- Curriculum reviewed against current hiring data, not run unchanged year after year
- Deployment and generative AI taught as core skills from day one, not a bonus module at the end
- Direct access for questions throughout the programme – no hand-off to a teaching assistant

Built for what comes **after** the notebook

Most data science courses stop at a trained model in a notebook. This one doesn't. **The 2026 market divides sharply between candidates who can build a model and candidates who can also ship one** – and that gap is the single biggest design constraint behind this curriculum.

That's why deployment and MLOps get a full two-week phase instead of a closing lecture, and why generative AI – RAG, LangChain, vector databases – is a core phase rather than an elective bolt-on at the end.

"Employers are prioritising candidates who can demonstrate real skills through practical projects and hands-on experience – not just theory."

The result mirrors how a working data scientist's year actually goes: statistics and Python first, classical machine learning next, then deep learning, generative AI, and finally the deployment skills that turn a model into something a business can actually run.

Who this is for

Working professionals moving from analytics, software or engineering into data science, and data analysts ready to go deeper into machine learning, deep learning and applied generative AI.

WHAT YOU WILL BE ABLE TO DO

- Build and evaluate **classical ML models** – regression, classification, ensembles – end to end
- Engineer features and explain model decisions using **SHAP**
- Design and train **deep learning models** in TensorFlow and PyTorch
- Build a **RAG application** using LangChain and a vector database
- Deploy a model as a **live, monitored API** using Docker and MLflow
- Use AI copilots as a working part of the data science workflow
- Show a deployed, end-to-end **capstone project** – chosen from eight real-world tracks – in interviews

Nineteen weeks, **nine** phases

01

Python & Statistics for Data Science

WEEKS 1-3

Advanced pandas and NumPy, probability, hypothesis testing, and the linear algebra that underpins every model later.

Python

NumPy / pandas

Statistics

OUTCOME Manipulate and reason about data like a working data scientist

02

SQL for Data Science

WEEK 4

A focused refresher on analytical SQL for sourcing and joining the features a model will actually use.

SQL

OUTCOME Pull and shape training data straight from a live database

03

Machine Learning Essentials

WEEKS 5-7

Regression, classification, decision trees and ensemble methods, with proper train/test discipline and evaluation metrics throughout.

Scikit-learn

XGBoost

Model Evaluation

OUTCOME Build and honestly evaluate a classical ML model end to end

04

Feature Engineering & Optimization

WEEK 8

Feature engineering, hyperparameter tuning, handling imbalanced data, and explaining model decisions with SHAP and LIME.

Feature Engineering

SHAP

OUTCOME Turn a mediocre model into a strong one, and explain why it works

05

Deep Learning

WEEKS 9-11

Neural networks from first principles, CNNs, RNNs, and an introduction to Transformer architecture in TensorFlow and PyTorch.

TensorFlow

PyTorch

CNNs / RNNs

OUTCOME Design and train a neural network for an image or sequence problem

Phases **six** to nine

06

NLP & Generative AI

WEEKS 12-14

NLP fundamentals, embeddings, vector databases, LangChain, retrieval-augmented generation, prompt engineering and fine-tuning basics.

LLMs

LangChain

RAG

OUTCOME Build a working RAG application on top of a real document set

07

MLOps & Deployment

WEEKS 15-16

Packaging a model as an API with FastAPI, containerising it with Docker, tracking experiments in MLflow, and monitoring for drift once it's live.

FastAPI

Docker

MLflow

OUTCOME Ship a model as a live, monitored API instead of a notebook

08

AI-Assisted Data Science & Communication

WEEK 17

Using AI copilots to accelerate modelling and analysis, then building a Power BI dashboard and a stakeholder-ready recommendation from your results.

AI Copilots

Power BI

Presentation

OUTCOME Turn a technical result into a decision someone else can act on

09

Capstone Project

WEEKS 18-19

One dataset, start to finish, from a choice of eight real-world tracks: modelling, evaluation and a deployed API - the centrepiece of your portfolio.

End-to-end project

OUTCOME Walk into interviews with a deployed, portfolio-ready project

More **production** experience than most programmes give you

220

CODING ASSIGNMENTS

30

**MODULE-END
ASSIGNMENTS**

8

MINI PROJECTS

1

**DEPLOYED CAPSTONE
PROJECT**

CAREER SUPPORT, BUILT IN

One-on-One Mentoring

Direct, personal guidance from your mentor throughout the programme – not a shared help desk.

Mock Interviews

Practice rounds structured around how data science interviews actually run, with real feedback.

Resume Preparation

A resume built around your actual projects and deployed models, not a generic template.

Interview Preparation

Structured practice for both technical and HR rounds, so nothing on interview day is unfamiliar.

+ DSA

+ Aptitude

ADD-ON

Dedicated DSA (Data Structures & Algorithms) and Aptitude sessions are provided as part of interview preparation – useful for learners also targeting broader technical interview rounds.

Eleven skills, **one** workflow

PY Python pandas, NumPy, scikit-learn	STAT Statistics Hypothesis testing, probability	SQL SQL Analytical queries for features
ML Classical ML Scikit-learn, XGBoost	DL Deep Learning TensorFlow, PyTorch	RAG GenAI & RAG LangChain, vector databases
OPS MLOps Docker, MLflow, FastAPI	PBI Power BI Dashboards for stakeholder reporting	GIT Git & GitHub (Essentials) Version control fundamentals
DSA DSA Data structures & algorithms for interviews	APT Aptitude Quantitative & logical reasoning practice	

WHAT YOU WALK AWAY WITH

- A **deployed, portfolio-ready capstone project** across ML, deep learning and MLOps
- Confidence building **RAG applications** with LangChain and vector databases
- The ability to explain model decisions to a non-technical stakeholder
- A skill set matched to **what employers are hiring for right now**

CAPSTONE

Every learner finishes with a **deployed, monitored model** – not just a notebook. **Eight real-world tracks** to choose from, overleaf.

A project at **every** milestone

MILESTONE PROJECTS ALONG THE WAY

WEEK 3	Statistical analysis & hypothesis testing project Apply descriptive statistics and hypothesis testing to a real dataset to test a business assumption.
WEEK 7	End-to-end machine learning model Build, evaluate and compare classical ML models on a real prediction problem.
WEEK 8	Model explainability project Use SHAP to explain what's actually driving a model's predictions.
WEEK 11	Deep learning classifier Train a CNN or RNN on a real image or sequence dataset.
WEEK 14	RAG-based Q&A application Build a working retrieval-augmented generation app using LangChain and a vector database.
WEEK 16	Deployed model API Package, containerise and deploy a model as a live, monitored API.

CHOOSE YOUR CAPSTONE TRACK - WEEKS 18-19

At Unicus, every learner has **the freedom to choose a capstone project** aligned to their own domain and interests, and **we actively encourage learners to pursue it** - the tracks below are a starting point, not a limit.

TRACK 01 E-commerce dynamic pricing Recommend the optimal price for a product using demand, competitor and seasonal signals. Python · Regression · Time Series	TRACK 02 Short-term rental price optimization Predict the optimal nightly price for a listing from property and market signals. Python · Regression · Time Series
TRACK 03 Content engagement prediction Predict which video titles and posting patterns drive the most viewer engagement. Python · NLP · Regression	TRACK 04 IT service ticket classification Automatically classify and route enterprise support tickets to the right team. Python · NLP · Classification
TRACK 05 Customer feedback & sentiment mining Mine reviews and tickets to surface themes and sentiment. NLP · Transformers · Python	TRACK 06 RAG-based domain chatbot Build a chatbot that answers questions grounded in real documents. LangChain · RAG · Vector DB
TRACK 07 Shadow API security risk detection Detect undocumented API calls and assign a risk score. Python · Anomaly Detection · MLOps	TRACK 08 Geographic risk clustering & mapping Cluster incident data into risk zones on an interactive map. Python · Clustering · Visualization

PROGRAMME FORMAT

Live & Mentor-Led

19 weeks - Mon to Fri - 1.5 hrs/day

Enrol now

7013265897 | enquiry@unicuslearning.com
Hyderabad, India | www.unicuslearning.com