AI for Data Science

Course Overview

Modern organizations rely on data systems that do far more than store information. They must support analytics, forecasting, visualization and, increasingly, AI-driven decision-making. This course prepares learners to use AI responsibly and effectively across the full data workflow—from querying and exploration, through visualization and modelling, to production-ready AI systems and agents.

Rather than focusing on theory alone, this course emphasizes applied, job-ready skills. Learners work with real tools, realistic business scenarios, and AI-assisted workflows to retrieve data, generate insights, build predictive models, and design reliable AI-enabled systems. Throughout the course, learners examine both the power and limitations of AI, then understand how to validate outputs, manage risk, and communicate results clearly to stakeholders.

By the end of the course, learners will have demonstrated the ability to design and operate end-to-end data and AI workflows suitable for real organizational environments.

Who This Course Is For

This course is designed for learners who want practical experience applying AI across data and analytics workflows, including:

- Data analysts and business analysts
- Data engineers and analytics engineers
- Technical professionals working with dashboards, reports, or forecasting
- Practitioners transitioning into AI-enabled roles
- Professionals responsible for deploying or evaluating AI systems

A basic familiarity with data concepts (e.g., datasets, queries, charts) is helpful. The course is structured to guide learners step-by-step from foundational tasks to more advanced, production-oriented workflows.

What You Will Learn

By completing this course, you will be able to:

- Translate business questions into effective, efficient data queries
- Use AI tools to assist with data exploration while validating outputs
- Design clear, decision-ready visualizations and dashboards
- Communicate insights using data storytelling best practices
- Build and refine predictive models with AI assistance
- Explain assumptions, uncertainty, and limitations in forecasts
- Design and operate AI agents within structured data workflows
- Apply guardrails, cost controls, and evaluation strategies

- Implement context, memory, retrieval, and evaluation for reliable AI systems
- Diagnose and debug common AI and data workflow failures

How the Course Works

- **Duration:** 8 weeks (24 total instructional hours)
- **Format:** Instructor-led sessions with guided labs
- **Learning Style:** Applied, skills-based, and project-focused

Each week includes:

- A focused lecture introducing core concepts
- A hands-on lab using industry-standard tools
- A short knowledge check or take-home assignment

From **Week 4 through Week 8**, learners work on a cumulative **Capstone Project** that integrates skills from across the course.

Weekly Course Outline

Week	Topics	Overview	Focus Areas
1	AI-Assisted Data Querying & Exploration (Foundations)	Learn how AI can support data querying and exploration while understanding its limitations. Practice translating business questions into queries and validating AI-generated outputs.	• Query formulation • Key Performance Indicators (KPIs) and business questions • Validating AI-generated queries • Avoiding hallucinated fields and inefficient queries
2	Efficient Query Design, Debugging & Optimization	Go deeper into performance, cost awareness, and reliability when querying data. Learn how to debug incorrect or inefficient AI-assisted queries.	• Query optimization • Debugging logic errors • Cost-aware querying • Performance trade-offs
3	Visualization, Reporting & Insight Communication	Learn how to transform queried data into clear, meaningful visuals. Practice selecting appropriate charts, designing dashboards, and communicating insights effectively.	• Visualization best practices • Chart selection • Dashboard design • Avoiding misleading or cluttered visuals
4	Predictive Modelling & Forecasting with AI Assistance	Introduce predictive modelling concepts and explore how AI can assist with forecasting and scoring. Begin the Capstone Project by defining a predictive or analytical use case.	• Feature selection • Time-series fundamentals • Scenario planning • Communicating uncertainty and assumptions
5	Model Refinement,	Iterate on predictive models and forecasts. Learn how to evaluate model quality, avoid	• Iterative model refinement • Evaluation metrics

Week	Topics	Overview	Focus Areas
	Evaluation & Stakeholder Alignment	overfitting, and ensure alignment with stakeholder needs.	• Explaining model limitations • Stakeholder communication
6	Designing & Operating AI Agents in Data Workflows	Learn how AI agents can be used to automate multi-step data workflows. Explore agent patterns, tool calling, and containerized execution.	• AI agent design patterns • Workflow automation • Containerization concepts • Debugging agent behaviour
7	Guardrails, Cost Control & Evaluation for AI Systems	Focus on reliability, safety, and operational concerns. Learn how to implement guardrails, monitor costs, and evaluate AI systems before deployment.	• Guardrails and safety checks • Cost monitoring and control • Evaluation strategies • Preventing silent failures and regressions
8	Context, Memory, Retrieval & Capstone Completion	Explore how context, memory, and retrieval affect AI behaviour in production systems. Complete and present the Capstone Project.	• Context windows and Retrieval-Augmented Generation (RAG) strategies • Memory management • Observability and tracing • End-to-end system reliability

Capstone Project: Applied Data & AI Workflow

The Capstone Project is the core applied component of the course. Learners design and demonstrate an end-to-end workflow that:

- Retrieves and validates data using AI-assisted querying
- Produces clear visualizations or reports aligned to a business question
- Incorporates predictive modelling or forecasting
- Uses AI tools or agents responsibly within a structured workflow
- Applies evaluation criteria to ensure reliability, cost awareness, and quality

The emphasis is not on building complex custom AI models, but on demonstrating that data and AI workflows work reliably, transparently, and appropriately for decision-making.

The capstone is designed to be:

- Practical and job-relevant
- Portfolio-ready
- A clear demonstration of applied AI-enabled data skills

Assessment & Completion

This course uses a **Complete / Incomplete** model focused on demonstrated skills rather than exams.

To successfully complete the course, learners must:

- Participate in weekly labs and activities
- Complete required assignments and knowledge checks
- Contribute to and present the Capstone Project

Completion indicates readiness to work with real-world data engineering systems.

What You'll Leave With

By the end of the course, you will have:

- Hands-on experience applying AI across data workflows
- A structured understanding of AI strengths and limitations
- An end-to-end applied project you can explain and defend
- Practical experience suitable for professional or portfolio use