JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

Advanced AI and Machine Learning Training Course for Professionals

jupiterprime.uk

info@jupiterprime.uk

Registered in the United Kingdom and Saudi Arabia

Advanced AI and Machine Learning Training Course for Professionals

AI and Machine Learning for professionals is a 2-week advanced training course for managers, analysts and technical leads who sponsor, specify or oversee model development, and who leave with a machine learning project charter for a live use case. Machine learning now sits inside forecasting, inspection, service and risk decisions, so organisations need people who can judge models rather than only commission them. Jupiter Prime Training runs this course on AI and machine learning for professionals in that position.

Why AI and Machine Learning Matters Now

Three forces shape AI and machine learning today: foundation models and large language models that put prediction and generation within reach of every function, cheaper cloud compute that shortens the path from prototype to production, and growing demands from boards, customers and auditors to explain how a model reaches its output. Organisations that handle this well choose the right problem, data and algorithm and keep models accurate after release; those that handle it badly fund projects that never leave the notebook or ship models that quietly drift. Examples come from retail demand forecasting and manufacturing quality inspection. The course rests on ISO/IEC 42001 for AI management systems and the NIST AI Risk Management Framework AI RMF.

Who Gains Most from AI and Machine Learning

- Product and business owners who frame prediction problems and accept or reject model results
- Analytics leads responsible for datasets, features and evaluation baselines
- Technical managers who oversee engineers building, deploying and monitoring models
- Risk and assurance specialists who challenge model behaviour, bias and documentation
- Digital transformation managers who choose between building, buying or fine-tuning models

The greatest gain comes when a business owner and the analytics and engineering leads for the same use case attend together and leave with one shared charter.

AI and Machine Learning Capabilities

Participants move from algorithm names to sound choices about problem framing, data, evaluation and deployment for real prediction and generation tasks.

- Specify supervised, unsupervised or generative approaches that match a business problem — use-case sponsors
- Assess model evaluation results using metrics suited to class balance and cost of error — model acceptance
- Weigh feature engineering and data labelling options against effort and accuracy gains — analytics leads
- Compare deep learning, natural language processing and computer vision architectures for given inputs — technical managers
- Design MLOps controls for deployment, monitoring and retraining — operations and assurance owners
- Build a machine learning project charter with scope, data, metrics and responsible AI checks — the sponsoring executive

Key AI and Machine Learning Terms

- Supervised Learning — training a model on labelled examples so it can predict a known target, such as demand or default, for new cases.
- Unsupervised Learning — finding structure in unlabelled data, such as customer clusters or anomalies, without a predefined target.
- Feature Engineering — creating and selecting input variables from raw data so a model can detect the patterns that matter.
- Overfitting — when a model learns noise in training data and performs well there but poorly on new data.
- Large Language Model LLM — a neural network trained on large text corpora to generate, summarise and classify language.
- MLOps — practices and tooling that version, deploy, monitor and retrain machine learning models reliably in production.

AI and Machine Learning Day by Day

Day 1: The Machine Learning Landscape

- AI, Machine Learning and Deep Learning Taxonomy — Choosing the Right Family
- Foundation Models and Compute Economics — Build Versus Fine-Tune Choices
- CRISP-DM Life Cycle — Framing Problems Before Selecting Algorithms
- Prediction Problem Types — Classification, Regression, Ranking and Generation
- Machine Learning Use-Case Canvas — Value, Feasibility and Data Availability

Day 2: Supervised Learning Methods

- Linear and Logistic Regression — Interpretable Baselines for Business Targets
- Decision Trees, Random Forests and Gradient Boosting — Tabular Accuracy
- Support Vector Machines and Nearest Neighbours — Small-Sample Problems
- Time-Series Forecasting Models — Demand and Capacity Planning Inputs
- Class Imbalance Techniques — Rare-Event Detection in Fraud and Faults

Day 3: Unsupervised Learning and Feature Engineering

- K-Means and Hierarchical Clustering — Customer and Asset Segmentation
- Principal Component Analysis — Dimensionality Reduction for Wide Datasets
- Isolation Forest and Autoencoder Anomaly Detection — Exception Monitoring
- Feature Encoding, Scaling and Interaction Terms — Input Quality Decisions
- Data Labelling Strategies — Cost, Consistency and Active Learning

Day 4: Model Evaluation and Selection

- Train, Validation and Test Splits — Leakage Prevention Rules
- Precision, Recall, F1 and ROC-AUC — Metric Choice by Cost of Error
- Cross-Validation and Hyperparameter Tuning — Model Selection Evidence
- Bias-Variance Trade-Off and Regularisation — Overfitting Controls
- Baseline and Champion-Challenger Comparison — Acceptance Criteria

Day 5: Deep Learning and Neural Networks

- Feedforward Networks and Backpropagation — When Depth Adds Accuracy
- Convolutional and Recurrent Architectures — Matching Networks to Input Types
- Transfer Learning From Pretrained Models — Reducing Data and Compute Needs
- Training Infrastructure and GPU Budgeting — Experiment Throughput Choices
- Neural Network Interpretability Methods — Saliency and Attribution Evidence

Day 6: Natural Language Processing

- Tokenisation and Embeddings — Representing Text for Models
- Text Classification and Sentiment Models — Routing and Voice-of-Customer Insight
- Named Entity Recognition — Extracting Fields From Documents
- Transformer Architecture and Attention — Why Context Length Matters
- Multilingual and Arabic Language Models — Coverage and Quality Testing

Day 7: Computer Vision

- Image Classification Pipelines — Defect and Document Recognition
- Object Detection and Segmentation — Inspection and Safety Monitoring
- Optical Character Recognition With Deep Learning — Paper-to-Data Workflows
- Data Augmentation and Synthetic Images — Small Dataset Strategies
- Edge Deployment of Vision Models — Latency and Hardware Constraints

Day 8: Generative AI and Large Language Models

- Prompt Engineering Patterns — Reliable Output for Business Tasks
- Retrieval-Augmented Generation — Grounding Answers in Organisational Content
- Fine-Tuning and Parameter-Efficient Adaptation — Domain Accuracy Choices
- LLM Evaluation and Hallucination Testing — Release Readiness Checks
- AI Agents and Tool Use — Automation Scope and Human Approval Points

Day 9: MLOps and Responsible AI

- Model Versioning, Registries and CI/CD Pipelines — Controlled Releases
- Data and Concept Drift Monitoring — Retraining Triggers
- Fairness Metrics and Bias Mitigation — Protected Group Outcomes
- Model Cards and Datasheets — Documentation for Reviewers
- ISO/IEC 42001 and NIST AI RMF Controls — Model Risk Register Entries

Day 10: Practising the Machine Learning Project Charter

- Use-Case Selection and Problem Statement — Applying the Use-Case Canvas
- Algorithm and Architecture Shortlist — Linking Data Types to Model Families
- Evaluation Metrics and Acceptance Thresholds — Business Sign-Off Criteria
- Deployment, Monitoring and Fairness Checks — Operating Controls
- Machine Learning Project Charter — Presentation to the Sponsoring Executive

Applying AI and Machine Learning After the Course

- Run a baseline-versus-gradient-boosting comparison on one existing prediction task and report precision, recall and cost of error to the business owner
- Audit one production model for data drift and missing model card documentation, and set retraining triggers with the engineering lead

AI and Machine Learning Questions FAQ

What is AI and machine learning and why does it matter now?

AI and machine learning are methods that let software learn patterns from data to predict, classify or generate outputs. They matter now because foundation models and cheaper compute put them within reach of every function, while scrutiny of model behaviour keeps rising.

Which roles gain most from AI and machine learning training for professionals?

Product owners, analytics leads, technical managers and risk specialists gain most. They frame problems, prepare data, oversee engineers or challenge models, so they need enough depth to judge algorithms, metrics and deployment controls.

How does advanced AI and machine learning training differ from an AI for business course?

It goes inside the models rather than staying at strategy level. You compare algorithms, read evaluation metrics, examine neural architectures and MLOps controls, at a depth suited to professionals who specify or oversee machine learning work.

What do you bring back from AI and machine learning training?

You bring back a machine learning project charter for a live use case. It sets the problem statement, data sources, candidate algorithms, evaluation thresholds, deployment controls and responsible AI checks your sponsor needs to approve the work.

AI and Machine Learning: The Way Forward

The machine learning project charter first informs whether a proposed use case receives data access and engineering capacity. Sponsors, analytics leads and engineering managers use it to agree scope and acceptance thresholds. After the first model release, revisit the evaluation metrics and drift signals, since production data often changes which algorithm and features deserve the next iteration.