
Machine Learning Model Development and Deployment Training Course

28 June – 2 July 2027
New York





Machine Learning Model Development and Deployment Training Course

Ref: 45_12023 **Date:** 28 June – 2 July 2027 **Location:** New York **Fees:** 7900 GBP

Programme Summary

Machine learning model development and deployment is a one-week training course for technical leads at advanced practitioner level, culminating in a production readiness dossier for one model from their own portfolio. Many organisations can build a promising prototype yet struggle to validate it, release it safely and keep it accurate once live data begins to shift. The programme follows the full lifecycle, from problem framing to monitoring and retraining, through hands-on work with realistic datasets. Mars Training Centre UK delivers this programme in machine learning model development and deployment.

The Business Case

Models that never leave the notebook return nothing on the investment in data and talent, while models released without validation or monitoring can quietly degrade, embed bias or fail an audit. Teams that apply a disciplined lifecycle move good models into service sooner, detect decline before it reaches customers and can explain each decision to risk and compliance reviewers.

- Time to Production: Shortens when experiments, packaging and release follow a repeatable pipeline.
- Model Risk Exposure: Falls through structured validation, bias testing and documented approval.
- Sustained Model Accuracy: Holds when drift is monitored and retraining is triggered by agreed thresholds.

Principles and Evidence Base

- CRISP-DM: Structures problem framing, data understanding and evaluation checkpoints in Sessions 1 and 2.
- ISO/IEC 23053:2022 Framework for AI Systems Using Machine Learning: Describes the components and lifecycle stages used to design the pipeline in Sessions 3 and 4.
- ISO/IEC 42001:2023 AI Management System: Frames responsible AI controls, roles and documentation in Session 4.
- Statistical Learning Theory: Underpins the bias-variance trade-off, cross-validation and model selection in Session 2.

Intended Participants

- Leads responsible for turning business questions into machine learning use cases.
- Practitioners who supervise model building, evaluation and code review within a data science team.
- Technical managers accountable for releasing and operating models in production systems.
- Analytics heads who must evidence model quality to risk, audit and business owners.



The programme serves banking and insurance, telecommunications, energy and utilities, and government digital services.

Learning Outcomes

1. Participants will be able to translate a business problem into a measurable machine learning task with success criteria.
2. Participants will be able to construct feature pipelines that prepare data consistently for training and inference.
3. Participants will be able to evaluate supervised and unsupervised models with metrics suited to the decision at stake.
4. Participants will be able to diagnose overfitting, leakage and bias and justify the selected model.
5. Participants will be able to design a deployment and monitoring plan covering drift detection and retraining triggers.
6. Participants will be able to compile model documentation that supports responsible AI review.

Programme Structure

Day 1: Framing Machine Learning Problems and Preparing Data

Which business problems suit machine learning, and is the data fit for them?

- CRISP-DM Business and Data Understanding Phases
- Target Definition, Success Metrics and Baselines
- Data Quality Profiling and Missing Value Treatment
- Feature Engineering: Encoding, Scaling and Aggregation
- Data Leakage Detection and Train, Validation and Test Splits

Day 2: Training and Evaluating Models

How are candidate models trained and compared fairly?

- Supervised Learning: Regression, Tree Ensembles and Gradient Boosting
- Unsupervised Learning: Clustering and Anomaly Detection
- Cross-Validation and Hyperparameter Tuning
- Evaluation Metrics: Precision, Recall, ROC-AUC and RMSE
- Experiment Tracking and Reproducibility

Day 3: Model Validation, Explainability and Bias Testing

Can the chosen model be trusted and explained before release?



- Independent Validation and Challenger Models
- Explainability With SHAP Values
- Fairness Metrics and Bias Mitigation
- Stress Testing on Edge Cases and Shifted Data
- Model Cards and Validation Reports

Day 4: Deployment and MLOps Monitoring

How does a validated model reach production and stay accurate?

- ISO/IEC 23053:2022 Lifecycle and Pipeline Components
- Packaging, Batch and Real-Time Serving Patterns
- CI/CD for Machine Learning and Model Registries
- Data and Concept Drift Monitoring
- ISO/IEC 42001:2023 Controls and Retraining Governance

Day 5: Applying Machine Learning Model Deployment

How do the lifecycle steps combine for a participant's own model?

- Framing a Participant Use Case With Success Criteria
- Reviewing the Feature Pipeline and Evaluation Evidence
- Setting Drift Thresholds and Retraining Triggers
- Peer Challenge of Validation and Bias Findings
- Production Readiness Dossier Presentation

Methodology

The programme is run as a data workshop. Participants work in small teams on anonymised datasets from credit risk, customer churn and equipment sensor cases, building, validating and deploying models in a shared environment. Facilitators act as risk reviewers, challenging each team's evidence before a model is allowed to progress to the next stage.

- Suggested activity: Finding the leakage that inflates a churn model's accuracy.
- Suggested activity: Writing a model card for a credit decision model.
- Suggested activity: Responding to a simulated drift alert on a sensor model.

Capstone Output

Each participant completes a production readiness dossier for one model. It records the problem framing and success criteria, the feature pipeline, evaluation and validation evidence, bias and explainability findings, the deployment pattern, drift thresholds and the retraining plan. It is used as the release pack submitted to model risk and business owners.



Frequently Asked Questions

Is machine learning model development and deployment suitable for experienced data practitioners?

Yes. It is designed for technical leads and practitioners who already write code for data analysis and know basic statistics. It assumes working familiarity with a data science language but no prior production deployment experience.

Which standards underpin machine learning model development and deployment in this programme?

It rests on the CRISP-DM process model, ISO/IEC 23053:2022 for machine learning system lifecycle, ISO/IEC 42001:2023 for AI management controls and statistical learning theory for model selection.

How is the machine learning production readiness dossier used at work?

You submit it as the release pack for a model, giving risk, audit and business owners the evidence they need to approve deployment and the monitoring plan to keep it accurate.

How does machine learning model deployment differ from an executive AI strategy programme?

It is hands-on and technical, focused on building, validating and operating models, whereas executive AI programmes address investment, governance and adoption decisions at board level.