JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

Machine Learning Literacy and Model Limitations for Non-Technical Leaders Training Course

jupiterprime.uk

info@jupiterprime.uk

Registered in the United Kingdom and Saudi Arabia

Machine Learning Literacy and Model Limitations for Non-Technical Leaders Training Course

Machine Learning Literacy and Model Limitations for Non-Technical Leaders is a 1-week Foundation training course for managers and professionals new to AI, in which each participant builds a Model Limitations Briefing Card for their function. As learned systems shape hiring, triage and customer service, leaders who cannot question a model's training evidence or error patterns inherit risks they never see. Jupiter Prime Training offers this course for leaders without a technical background.

Why Machine Learning Literacy Matters Now

Three forces shape how leaders meet machine learning today: generative assistants arriving inside everyday software, stakeholders and auditors asking who answers for automated outcomes, and model errors that fall unevenly on customers and staff. Leaders who can read a confusion matrix and spot drift catch failures early; those who treat outputs as facts find bias through complaints. The course draws on ISO/IEC 42001 Artificial Intelligence Management System and the NIST AI Risk Management Framework AI RMF for its oversight vocabulary.

Who Gains Most from Machine Learning Literacy

- Functional leaders whose teams now rely on predictions, rankings or generated text in daily decisions
- Operations heads responsible for where automated recommendations enter a process and who reviews them
- People and customer leaders accountable for fair treatment when a model scores individuals
- Risk, audit and quality leaders who must ask informed questions of data science teams
- Product and service owners who brief executives on what an AI feature can and cannot do

Cross-functional leadership groups gain most, sharing one vocabulary for error, bias and drift.

Machine Learning Literacy Capabilities

Participants leave able to question how a model learned, where its errors fall and what its confidence really means.

- Assess whether a training sample represents the people and cases a model will meet — fairness reviews
- Weigh false positives against false negatives for a given use — acceptable error decisions
- Test vendor accuracy claims against benchmark context and held-out evidence — procurement conversations
- Specify human review checkpoints where model errors carry the highest cost — process owners

- Communicate known model limitations in plain language — executives and frontline staff
- Apply drift and calibration warning signs to decide when a model needs revisiting — data science teams

Key Machine Learning Literacy Terms

- Training data — the labelled or unlabelled examples a model studies to find patterns it later applies to new cases.
- Overfitting — when a model memorises quirks of its training examples and performs poorly on genuinely new cases.
- Confusion matrix — a table counting correct and incorrect predictions by type, revealing which errors a model makes most.
- Model drift — gradual loss of accuracy as real-world conditions move away from the data the model learned from.
- Algorithmic bias — systematic differences in model errors or outcomes between groups, often inherited from skewed training data.
- Hallucination — fluent but false or invented content produced by a generative model with apparent confidence.

Machine Learning Literacy Day by Day

Day 1: How Machines Learn from Examples

- Supervised, Unsupervised and Reinforcement Learning — Matching Problem Types to Approaches
- Features, Labels and Training Examples — What a Model Actually Sees
- Training, Validation and Test Splits — Why Held-Out Evidence Matters
- Rules-Based Software Versus Learned Behaviour — Setting Leader Expectations
- Large Language Models and Next-Token Prediction — Demystifying Generative Output

Day 2: Data Quality and Representativeness

- Sampling Gaps and Under-Represented Groups — Who the Model Ignores
- Label Noise and Annotation Consistency — Trusting the Ground Truth
- Data Drift and Concept Drift — When Yesterday's Patterns Expire
- Proxy Variables and Leakage — Hidden Shortcuts Behind Impressive Scores
- Datasheets for Datasets — Recording Provenance for Reviewers

Day 3: Where Models Fail and Show Bias

- Overfitting and Underfitting — Reading Learning Curves Without Code
- Precision, Recall and the Confusion Matrix — Choosing Acceptable Error Types
- Calibration and Confidence Scores — When Probabilities Mislead Decision Makers
- Algorithmic Bias and Fairness Metrics — Weighing Disparate Error Rates
- Hallucination and Fabricated Citations — Limits of Generative Text

Day 4: Questioning Models and Their Claims

- Model Cards — Interrogating Intended Use and Known Limitations
- Explainability Techniques Such as Feature Importance — Questions for Data Scientists
- Vendor Accuracy Claims and Benchmark Context — Spotting Misleading Comparisons
- NIST AI Risk Management Framework AI RMF — Mapping and Measuring Model Risk
- ISO/IEC 42001 — Roles and Oversight in an AI Management System

Day 5: Applying Model Literacy in Practice

- Task Suitability Grid — Where Prediction Helps and Where Judgement Stays Human
- Human-in-the-Loop Checkpoints — Placing Review Where Errors Cost Most
- Error Case Walkthrough — Applying Bias and Drift Lessons to a Live Example
- Plain-Language Limitation Statements — Briefing Executives and Frontline Staff
- Model Limitations Briefing Card — Finalising the Shared Reference

Applying Machine Learning Literacy After the Course

- Run a confusion matrix walkthrough with the data science team on one live model, agreeing which error type the function can tolerate.
- Request or draft a model card for an AI feature already in use, recording intended use, training data sources and known blind spots.

Machine Learning Literacy Questions FAQ

What is machine learning literacy for non-technical leaders and why does it matter now?

It is the ability to question how a model learned and where it errs. It matters because learned systems already shape everyday decisions, and leaders stay accountable.

Which roles gain most from machine learning literacy and model limitations training?

Managers and professionals without a technical background gain most. That includes functional, risk and product leaders who rely on model outputs.

How does machine learning literacy differ from AI strategy or data science courses?

It sits between them. Strategy courses focus on direction and portfolio choices, data science courses on building models; you instead learn to interrogate models, their training evidence and their failure patterns without writing code.

What do you bring back from machine learning literacy and model limitations training?

You bring back a Model Limitations Briefing Card for your function. It sets out which errors matter, warning signs of drift and bias, and questions to ask before trusting an output.

Machine Learning Literacy: The Way Forward

The Model Limitations Briefing Card first informs where a human must confirm a model's output before it reaches a customer or employee. Leaders and reviewers use it when an AI feature arrives or misbehaves. After its first use, revisit the acceptable error choices, since what a function tolerates often shifts once real cases accumulate.