JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

Tabular Prediction Models Hands-On Course for Business Analysts

21 – 25 November 2027
Muscat

Tabular Prediction Models Hands-On Course for Business Analysts

Ref: 1351_50604 **Date:** 21 – 25 November 2027 **Location:** Muscat **Fees:** 6500 GBP

Tabular Prediction Models for business analysts is a 1-week hands-on training course for analysts, project staff and junior technical professionals who want to move from theory to a working model, and it produces a Prediction Model Evaluation Report on one business dataset. Most organisational records sit in rows and columns, and the ability to turn them into tested forecasts is now expected well beyond specialist data science units. Jupiter Prime Training runs this course on Tabular Prediction Models.

Why Tabular Prediction Models Matter Now

Three forces push this skill towards analysts. Low-code platforms and open-source libraries have removed many technical barriers to fitting a model. Business owners in sales, finance and operations now ask for churn, demand and risk scores rather than static reports. Auditors and risk functions, meanwhile, want evidence that a score was validated before it influences customers or budgets. Analysts who can frame, fit and honestly test a model shorten the path from question to decision; those who cannot produce numbers nobody trusts. The course refers to the NIST AI Risk Management Framework AI RMF and ISO/IEC 42001 for documentation and validation expectations.

Who Gains Most from Tabular Prediction Models

- Analysts who prepare monthly reports and are now asked for forecasts
- Project staff who scope and test prediction features in business applications
- Junior developers who integrate model scores into dashboards or workflows
- Finance and operations staff who maintain spreadsheet-based estimates
- Quality and risk staff who review model outputs before release

An analytics or business intelligence unit gains most when two or three of its analysts attend together and return with a shared modelling routine.

Tabular Prediction Models Capabilities

Participants leave able to carry a prediction question from a vague request to a validated, explained result on real records.

- Frame a business question as a regression or classification target — requesting managers
- Build features from transactional and customer records — data owners
- Test models with hold-out sets and cross-validation — reviewers and auditors
- Compare a model against a simple baseline and a cost of error — budget holders
- Explain which inputs drive predictions using feature importance — business users
- Specify monitoring checks before a score goes live — operations teams

Key Tabular Prediction Models Terms

- Target variable — the column a model learns to predict, such as next month's sales or whether a customer leaves.
- Feature engineering — creating useful input columns from raw records, for example days since last purchase.
- Baseline model — a simple rule or average against which a trained model must show improvement.
- Cross-validation — repeatedly training and testing on different slices of data to obtain a stable performance estimate.
- Class imbalance — when one outcome is much rarer than another, making plain accuracy misleading.
- Feature importance — a measure of how much each input column contributes to a model's predictions.

Tabular Prediction Models Day by Day

Day 1: Framing Prediction Problems

- Supervised Learning Foundations — Regression Versus Classification Choices
- Problem Framing Canvas — Turning Requests Into Measurable Targets
- Baseline Rules and Averages — Setting the Bar to Beat
- Cost of Error Matrix — Weighing False Alarms Against Misses
- Retail Churn and Utility Demand Cases — Two Contrasting Targets

Day 2: Preparing Records for Modelling

- Data Profiling in Spreadsheets and Notebooks — Spotting Gaps Early
- Missing Value Treatment — Imputation Versus Exclusion Decisions
- Feature Engineering From Transactions — Recency, Frequency and Value
- Categorical Encoding Methods — Handling Regions, Products and Channels
- Leakage Checks — Removing Information Unavailable at Prediction Time

Day 3: Fitting and Comparing Models

- Linear and Logistic Regression — Interpretable Starting Points
- Decision Trees and Random Forests — Capturing Non-Linear Patterns
- Gradient Boosting Overview — When Extra Accuracy Justifies Complexity
- Hyperparameter Tuning Basics — Avoiding Endless Trial and Error
- Model Comparison Table — Choosing Between Candidate Models

Day 4: Testing, Explaining and Monitoring

- Hold-Out and Cross-Validation Designs — Honest Performance Estimates
- Error Measures for Regression and Classification — Matching Business Impact
- Class Imbalance Remedies — Resampling and Threshold Adjustment
- Feature Importance and Partial Dependence — Explaining Drivers to Users
- NIST AI RMF Measure Function — Documenting Validation Evidence

Day 5: Practising the Prediction Model Evaluation Report

- Insurance Claims Use Case Framing — Applying the Problem Framing Canvas
- Feature Build and Leakage Review — Applying Preparation Methods End to End
- Candidate Model Bake-Off — Applying Comparison and Validation Designs
- Monitoring Checklist Drafting — Drift Alerts and Retraining Triggers
- Prediction Model Evaluation Report — Presenting Results to Decision Makers

Applying Tabular Prediction Models After the Course

- Rebuild one existing spreadsheet forecast as a baseline and test a trained model against it using cross-validation.
- Run a leakage and class imbalance check on a live dataset your unit plans to use for scoring.

Tabular Prediction Models Questions FAQ

What are Tabular Prediction Models and why do they matter now?

Tabular Prediction Models are machine learning models trained on rows and columns of business records to forecast values or classify cases. They matter now because analysts are expected to deliver tested forecasts, not just descriptive reports.

Which roles gain most from Tabular Prediction Models training?

Reporting analysts, project staff, junior developers, finance and operations estimators, and quality reviewers gain most, since each either builds, integrates or checks prediction scores.

How do Tabular Prediction Models differ from an advanced machine learning engineering course?

Tabular Prediction Models focus on proven methods for structured business data at practitioner depth. An engineering course covers deep learning, pipelines and deployment infrastructure; this course stays with framing, preparation, validation and explanation.

What do you bring back from Tabular Prediction Models training?

You bring back a Prediction Model Evaluation Report for one use case, with a baseline comparison, validation results and a monitoring checklist.

Tabular Prediction Models: The Way Forward

The Prediction Model Evaluation Report gives decision makers a clear record of what a model predicts, how well and under which limits. Its first use is usually a go or no-go decision on a churn, demand or claims score, shared by analysts and budget holders. Revisit it after the first months of live use, when drift and real error costs become visible.