JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

Advanced Applied AI for Professionals Training Course

26 April – 7 May 2027
Vienna

Advanced Applied AI for Professionals Training Course

Ref: 5_148 **Date:** 26 April – 7 May 2027 **Location:** Vienna **Fees:** 12400 GBP

Advanced Applied AI for Professionals is a 2-week training course for professionals and team leads who use AI in daily work, and it ends with an applied AI solution blueprint each participant takes back to a real workflow. Low-code platforms and generative models now let non-specialists build working AI solutions, so organisations need people who can frame problems, prepare data and judge model output with care. Jupiter Prime Training runs this course on applied AI for professionals.

Why Applied AI for Professionals Matters Now

Three forces shape hands-on AI work today: AutoML and low-code platforms that move model building out of specialist labs, line managers who expect faster answers from operational data, and information security and audit functions that demand traceable, explainable outputs. A hospital scheduling unit that classifies referrals with a well-tested model frees clinicians' time; a logistics firm that deploys an untested forecast can mis-plan fleets for months. Careful problem framing, evaluation and monitoring make the difference. The course rests on the Cross-Industry Standard Process for Data Mining CRISP-DM, ISO/IEC 42001 for AI management systems and the NIST AI Risk Management Framework AI RMF.

Who Gains Most from Applied AI for Professionals

- Those who turn recurring operational problems into prediction, classification or text-generation tasks
- Those who prepare, clean and label spreadsheets, records and documents for model building
- Those who configure low-code AI platforms, automations and prompt workflows for colleagues
- Those who test model accuracy and decide whether an AI output is fit to act on
- Those who keep deployed AI solutions running, monitored and documented

Operational units such as finance, customer operations and quality gain most when an analyst, a process owner and a technical colleague attend together and build solutions around one shared workflow.

Applied AI for Professionals Capabilities

Analysts, process owners and team leads leave able to carry an AI idea from a framed question to a monitored solution:

- Specify an AI problem statement with target variable, inputs and success threshold — for sponsors approving the build
- Build a cleaned, labelled and split dataset ready for training — for reliable model results
- Test classification, regression and clustering approaches on low-code platforms — for choosing the right technique
- Design prompt chains and retrieval-augmented generation workflows — for document-heavy tasks
- Calculate precision, recall, error rates and confusion matrices — for go or no-go decisions
- Apply monitoring, documentation and bias checks to deployed models — for audit and risk owners

Key Applied AI for Professionals Terms

- CRISP-DM — a six-phase process for data projects: business understanding, data understanding, preparation, modelling, evaluation and deployment.
- Feature engineering — creating and selecting the input variables a model uses, often the biggest driver of accuracy.
- AutoML — software that automatically tries algorithms and settings to find a well-performing model with little coding.
- Retrieval-augmented generation — a method that feeds a language model passages from approved documents so its answers rest on trusted sources.
- Confusion matrix — a table comparing predicted and actual classes, showing true and false positives and negatives.
- Overfitting — when a model memorises training data and performs poorly on new, unseen cases.

Applied AI for Professionals Day by Day

Day 1: The Applied AI Landscape and Lifecycle

- Machine Learning, Deep Learning and Generative Models — Choosing the Right Family
- CRISP-DM Phases — Structuring an Applied AI Project
- Low-Code and AutoML Platforms — Capabilities and Limits for Non-Specialists
- Human-in-the-Loop Design — Where People Check Model Output
- Personal Workflow Audit — Spotting Tasks Suited to AI

Day 2: Framing Problems for AI

- Problem Statement Canvas — Target, Inputs, Users and Thresholds
- Prediction Versus Classification Versus Generation — Matching Task to Technique
- Baseline Rules and Heuristics — Knowing What AI Must Beat
- Success Criteria and Error Costs — Agreeing Acceptable Mistakes With Sponsors
- Feasibility Screening — Data Volume, Labels and Access

Day 3: Data Preparation and Feature Engineering

- Data Profiling — Missing Values, Outliers and Duplicates
- Cleaning and Transformation Pipelines — Repeatable Preparation Steps
- Labelling Strategies — Manual, Rule-Based and Assisted Annotation
- Feature Engineering Techniques — Dates, Categories and Text Inputs
- Train, Validation and Test Splits — Avoiding Leakage Between Sets

Day 4: Machine Learning Methods in Practice

- Regression Models — Forecasting Quantities Such as Demand and Cost
- Classification Models — Decision Trees, Random Forests and Boosting
- Clustering — Segmenting Customers, Assets or Cases
- Time Series Forecasting — Seasonality and Trend Handling
- AutoML Experiments — Comparing Algorithms Without Code

Day 5: Generative AI and Prompt Engineering Workflows

- Prompt Patterns — Role, Context, Examples and Output Format
- Prompt Chaining — Breaking Complex Tasks Into Steps
- Retrieval-Augmented Generation — Grounding Answers in Approved Documents
- Structured Extraction — Turning Emails and Forms Into Records
- Hallucination Controls — Citations, Constraints and Human Checks

Day 6: Evaluating Model Performance

- Confusion Matrix, Precision and Recall — Weighing False Alarms Against Misses
- Error Metrics for Forecasts — MAE, RMSE and Percentage Error
- Cross-Validation — Checking Stability Across Data Samples
- Overfitting Diagnosis — Learning Curves and Simplification
- Generative Output Evaluation — Rubrics, Spot Checks and Test Sets

Day 7: Automation and Integration

- Workflow Automation Platforms — Triggering Models From Business Events
- APIs and Connectors — Linking AI Services to Existing Systems
- Intelligent Document Processing — Invoices, Claims and Applications
- AI Agents and Tool Use — Delegating Multi-Step Tasks Safely
- Exception Handling Design — Routing Uncertain Cases to People

Day 8: Deployment and Monitoring

- Deployment Options — Batch Scoring, Real-Time Services and Embedded Tools
- Model Drift Detection — Tracking Input and Accuracy Changes
- Version Control and Model Registers — Knowing What Runs Where
- Retraining Triggers — When and How to Refresh Models
- Operational Dashboards — Usage, Accuracy and Business Effect

Day 9: Responsible AI in Everyday Use

- ISO/IEC 42001 Controls — Roles, Records and Impact Assessment
- NIST AI RMF Practices — Mapping and Measuring Solution Risks
- Fairness and Bias Testing — Comparing Outcomes Across Groups
- Explainability Techniques — Feature Importance and Example-Based Reasons
- Data Privacy and Confidentiality — Safe Handling in AI Tools

Day 10: Building the Applied AI Solution

- Problem Canvas Refinement — Applying Day Two Framing to a Live Workflow
- Prototype Build — Applying Day Four or Day Five Methods
- Evaluation Report — Applying Day Six Metrics to the Prototype
- Monitoring and Control Checklist — Applying Day Eight and Nine Safeguards
- Applied AI Solution Blueprint Assembly — Design, Evidence and Rollout Steps

Applying Applied AI for Professionals After the Course

- Rebuild the course prototype on the unit's own data and present its confusion matrix and error costs to the process owner.
- Set up a drift and accuracy dashboard for one deployed model or prompt workflow and agree a retraining trigger with its users.

Applied AI for Professionals Questions FAQ

What is applied AI for professionals, and why does it matter now?

Applied AI for professionals is the hands-on practice of framing work problems, preparing data, building models or prompt workflows and checking their output. It matters now because low-code and generative tools put model building in non-specialists' hands, and errors reach decisions quickly.

Which roles gain most from advanced applied AI for professionals training?

Analysts, process owners, operations team leads and technical colleagues who automate workflows gain most from advanced applied AI for professionals training. They build or run AI solutions day to day and must judge when outputs can be trusted.

How is an advanced applied AI for professionals course different from an AI strategy course?

An advanced applied AI for professionals course is hands-on: participants frame, build, test and monitor solutions themselves. An AI strategy course works at portfolio and investment level, deciding which initiatives to fund rather than how to build them.

What do participants bring back from applied AI for professionals training?

Participants bring back a framed problem canvas, a prepared dataset, a tested prototype, an evaluation report and a monitoring checklist for one real workflow. You can use them straight away to move a working AI solution towards use.

Applied AI for Professionals: The Way Forward

The applied AI solution blueprint sets out one workflow's design, test evidence and rollout steps in a form colleagues can review. The first decision it informs is whether the prototype is accurate enough to run alongside current practice. Process owners and analysts use it together, and the error-cost thresholds are worth revisiting once real users have worked with the output.