JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

Advanced AI-Assisted Data Analysis Workflows Training Course

4 – 15 July 2027
Cairo

Advanced AI-Assisted Data Analysis Workflows Training Course

Ref: 292_9631 **Date:** 4 – 15 July 2027 **Location:** Cairo **Fees:** 6500 GBP

AI-Assisted Data Analysis Workflows is a 2-week advanced training course for analysts and analytics leads who use AI assistants, AutoML and natural-language querying in daily work, and each participant builds an AI-Assisted Analysis Workflow with validation checkpoints. Generative AI now drafts queries, cleans tables and writes summaries, so organisations need analysts who can prove those outputs are right. Jupiter Prime Training delivers this course for practising analysts.

Why AI-Assisted Data Analysis Workflows Matter Now

Three pressures shape this subject. Generative AI assistants have moved into spreadsheets, notebooks and query editors, so analysts receive machine-written code and narratives every hour. Business leaders expect faster answers, yet a single invented figure in a board summary can damage trust in the whole analytics function. Privacy and risk teams also ask where sensitive records travel when prompts leave the organisation. Analysts who pair AI speed with disciplined checking shorten turnaround while keeping figures defensible; those who do not publish confident errors. The course rests on the Cross-Industry Standard Process for Data Mining CRISP-DM, used to place AI assistance inside each analytical stage.

Who Gains Most from AI-Assisted Data Analysis Workflows

- Analysts who prepare, query and report on business data with AI assistants
- Analytics leads responsible for checking AI-produced figures before release
- Financial and operational analysts who monitor transactions or readings for anomalies
- Business intelligence specialists building natural-language querying for colleagues
- Data quality owners who must trace every machine-proposed transformation

An analytics function gains most when several analysts attend together and agree one shared prompt library and one set of validation gates.

AI-Assisted Data Analysis Workflows Capabilities

Participants learn to let AI draft, clean and model while they keep firm control of evidence, sign-off and disclosure.

- Design validation checkpoints between AI-assisted preparation, querying, modelling and reporting stages — analytics leads
- Build structured prompts and text-to-SQL routines grounded in approved metric definitions — report consumers
- Assess AutoML leaderboards for leakage, overfitting and unstable drivers — model owners
- Apply anomaly detection with tuned thresholds and triage routing — finance and operations owners
- Test AI summaries for hallucinated figures and unsupported causes — executive readers

- Specify data minimisation and disclosure rules for AI use — privacy and risk teams

Key AI-Assisted Data Analysis Workflows Terms

- Hallucination — a fluent AI output that states a number, fact or cause not supported by the underlying data.
- Text-to-SQL — the generation of database queries from plain-language questions by a language model, needing review before execution.
- Retrieval-augmented generation — a technique that feeds selected internal documents to a language model so answers cite approved sources.
- AutoML — software that automatically tries algorithms, features and settings, then ranks candidate models on a chosen metric.
- Anomaly detection — methods that flag records or periods departing sharply from expected patterns for human investigation.
- Validation checkpoint — a defined gate where an analyst verifies AI output against evidence before the next stage proceeds.

AI-Assisted Data Analysis Workflows Day by Day

Day 1: Where AI Enters the Analyst's Day

- Analyst Task Inventory — Deciding Which Steps AI Assistants Should Handle
- Generative AI Capability Limits — Weighing Fluency Against Factual Reliability
- CRISP-DM Business Understanding — Anchoring AI Use to a Stated Question
- Human-in-the-Loop Patterns — Choosing Where Analysts Must Approve Outputs
- AI Usage Policy Review — Checking Permitted Data and Tools

Day 2: AI-Assisted Data Preparation

- AI-Suggested Cleaning Rules — Accepting or Rejecting Proposed Transformations
- Automated Data Profiling — Spotting Missing Values, Outliers and Type Errors
- Fuzzy Matching and Entity Resolution — Merging Customer or Supplier Records
- Synthetic Test Rows — Exercising Pipelines Without Exposing Personal Records
- Preparation Audit Log — Recording Every Machine-Proposed Change

Day 3: Prompting for Analytical Work

- Structured Prompt Templates — Specifying Role, Context, Format and Constraints
- Few-Shot Examples for Tabular Tasks — Steering Consistent Column Handling
- Chain-of-Steps Decomposition — Breaking Complex Questions Into Checkable Parts
- Prompt Libraries and Versioning — Sharing Proven Prompts Across Analysts
- Prompt Injection Awareness — Protecting Workflows From Untrusted Source Text

Day 4: Natural-Language Querying

- Text-to-SQL Generation — Reviewing Joins, Filters and Aggregations Before Running
- Semantic Layer Grounding — Tying AI Queries to Approved Metric Definitions
- Retrieval-Augmented Generation Over Documents — Answering From Cited Internal Sources
- Query Result Reconciliation — Matching AI Answers to Known Control Totals
- Conversational Analytics Interfaces — Deciding Who May Self-Serve Answers

Day 5: AI Summarisation and Narrative

- Automated Insight Summaries — Checking Every Claim Against Underlying Figures
- Hallucination Detection Techniques — Flagging Unsupported Numbers and Invented Causes
- Executive Brief Drafting — Editing Machine Narratives for Tone and Accuracy
- Chart Captioning With AI — Keeping Descriptions Faithful to Axes
- Citation and Source Tagging — Linking Each Statement to Its Table

Day 6: Automated Machine Learning

- AutoML Leaderboards — Interpreting Ranked Candidates Beyond the Top Score
- Feature Engineering Suggestions — Screening Leakage and Proxy Variables
- CRISP-DM Modelling — Setting Baselines Before Trusting Automated Searches
- Cross-Validation Design — Matching Splits to Time Order and Groups
- Model Explanation Reports — Presenting Drivers to Non-Technical Owners

Day 7: Anomaly Detection

- Statistical Control Limits — Separating Ordinary Variation From Real Shifts
- Isolation Forest and Density Methods — Scoring Unusual Transactions or Readings
- Seasonal Decomposition for Time Series — Removing Expected Peaks Before Alerting
- Alert Threshold Tuning — Balancing Missed Incidents Against Alert Fatigue
- Anomaly Triage Workflow — Routing Flags to Finance or Operations Owners

Day 8: Validating AI-Generated Insights

- Validation Checkpoint Design — Placing Gates Between Each AI-Assisted Stage
- Ground-Truth Sampling — Hand-Checking a Fixed Share of AI Outputs
- Reproducibility Tests — Rerunning Prompts and Models for Stable Answers
- Bias and Fairness Checks — Comparing Results Across Customer or Staff Segments
- CRISP-DM Evaluation — Signing Off Findings Against the Original Question

Day 9: Risk, Privacy and Accountability

- Data Minimisation for AI Prompts — Removing Identifiers Before Submission
- Model and Vendor Risk Questions — Assessing Retention, Training Use and Hosting
- Decision Accountability Matrix — Naming Who Owns Each AI-Influenced Conclusion
- Disclosure Statements in Reports — Telling Readers Where AI Contributed
- Incident Handling for AI Errors — Correcting Published Figures Quickly

Day 10: Practising the AI-Assisted Analysis Workflow

- Retail Sales Case — Applying AI Preparation, Text-to-SQL and Summaries
- Hospital Operations Case — Applying Anomaly Detection and Triage Routing
- AutoML Churn Case — Applying Leaderboards, Explanations and Cross-Validation
- Peer Validation Exercise — Testing Checkpoints Against Planted AI Errors
- AI-Assisted Analysis Workflow — Completing Stages, Validation Gates and Owners

Applying AI-Assisted Data Analysis Workflows After the Course

- Introduce ground-truth sampling on one recurring AI-assisted report, hand-checking a fixed share of figures and logging every correction.
- Publish a shared prompt library with versioned text-to-SQL templates tied to approved metric definitions.

AI-Assisted Data Analysis Workflows Questions FAQ

What are AI-assisted data analysis workflows, and why do they matter now?

They are analytical processes in which AI tools prepare data, write queries, build models or draft summaries while analysts verify each output. They matter now because AI assistants are built into everyday analysis software, and unchecked machine output can reach decision makers within minutes.

Which roles gain most from AI-assisted data analysis workflows training?

Practising analysts, analytics leads, financial and operational analysts and data quality owners gain most. You benefit if you already use AI assistants for cleaning, querying or reporting in sectors such as retail, healthcare, banking or energy and must defend the figures.

How do AI-assisted data analysis workflows differ from a Power BI techniques course?

They are platform-neutral and centre on supervising AI output rather than mastering one reporting tool. A Power BI course builds models and DAX measures; this subject covers prompting, text-to-SQL, AutoML, anomaly detection and validation gates across whichever tools an analyst uses.

What do you bring back from AI-assisted data analysis workflows training?

You bring back a mapped analytical process for one real report, showing where AI assists, which checkpoint verifies each output, who signs off and how AI involvement is disclosed to readers.

AI-Assisted Data Analysis Workflows: The Way Forward

The AI-Assisted Analysis Workflow first decides which recurring report receives AI help and where the validation gates sit. Analysts, analytics leads and data quality owners use it to assign sign-off and correction duties. After its first cycle, the ground-truth sample rate and anomaly thresholds deserve another look, since real error rates show whether checking is too light or too heavy.