JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

Generative AI Hallucination Control Training Course for Professionals

22 – 26 November 2026
Cairo

jupiterprime.uk

info@jupiterprime.uk

Registered in the United Kingdom and Saudi Arabia

Generative AI Hallucination Control Training Course for Professionals

Ref: 1622_59381 **Date:** 22 – 26 November 2026 **Location:** Cairo **Fees:** 4000 GBP

Generative AI hallucination control is a 1-week training course for managers and professionals who rely on AI assistants for drafts, summaries and figures, and who leave with an Answer Verification Protocol for their function. Fluent but fabricated answers now reach reports, client letters and board papers, so checking machine-written claims before they are used has become a core professional habit. Jupiter Prime Training runs this course on generative AI hallucination control.

Why Generative AI Hallucination Control Matters Now

Three pressures shape generative AI hallucination control today: assistants embedded in everyday office software, executives who expect faster drafts, and clients or auditors who hold the organisation responsible for every stated fact. A team that checks invented citations, misquoted figures and confident guesses keeps its credibility; one that does not spreads errors that are slow and costly to retract. The course rests on ISO/IEC 42001, the international standard for an artificial intelligence management system, and the NIST AI Risk Management Framework NIST AI RMF, which treats validity and reliability as core trust characteristics.

Who Gains Most from Generative AI Hallucination Control

- Professionals who turn assistant drafts into reports, briefings and correspondence
- Managers who sign off documents where AI wrote part of the text
- Analysts who ask assistants for figures, comparisons and summaries of long files
- Quality and risk officers who set review rules for AI-written material
- Knowledge and records staff who maintain the source libraries assistants draw on

A department that sends a sign-off manager together with two or three frequent assistant users gains most, as drafters and approvers agree one routine.

Generative AI Hallucination Control Capabilities

Participants move from trusting fluent text to testing it, claim by claim, against evidence and risk.

- Assess why a language model produces a plausible but false statement — choice of tasks fit for assistants
- Apply claim extraction and source tracing to a generated draft — sign-off by the accountable manager
- Design prompts and grounding that narrow the room for invention — teams writing recurring documents
- Calculate a sampling rate for checking high-volume assistant output — quality and risk officers
- Specify human approval thresholds by consequence of error — department heads and auditors

- Build an error log that turns each caught fabrication into a lesson — owners of the source library

Key Generative AI Hallucination Control Terms

- Hallucination — a fluent, confident generated statement that is false or invented.
- Grounding — giving the model approved source passages so it answers from them.
- Claim extraction — breaking a draft into single checkable statements such as figures, names, dates and quotations.
- Fabricated citation — a reference or quotation the model presents as real that does not exist.
- Consequence tiering — ranking outputs by the harm an error would cause, so checking effort follows the stakes.
- Human-in-the-loop sign-off — a named person approving a generated output before it leaves the team or informs a decision.

Generative AI Hallucination Control Day by Day

Day 1: How Fluent Text Goes Wrong

- Next-Word Prediction Basics — Why Fluency Never Proves Accuracy
- Training Cut-Offs and Missing Context — Limits on Current Questions
- Text, Image and Code Failure Patterns — Matching Checks to Output Type
- Hallucination Case Files From Finance and Healthcare — Cost of Unchecked Claims
- Task Suitability Grid — Where Assistants Help and Where They Mislead

Day 2: Spotting Fabrication in Drafts

- Claim Extraction Technique — Splitting Drafts Into Checkable Statements
- Red-Flag Patterns — Precise Figures, Vague Sources and Odd Quotations
- Citation and Link Tracing — Confirming References Really Exist
- Numerical Cross-Checks — Reconciling Totals, Ratios and Units
- Code and Formula Spot Tests — Running Small Cases Before Reuse

Day 3: Reducing Invention at the Source

- Grounded Prompting With Approved Passages — Narrowing Answers to Evidence
- Retrieval From Curated Libraries — Choosing Which Documents Assistants See
- Instructions to Admit Uncertainty — Inviting Gaps Instead of Guesses
- Confidential Data Handling Rules — What Never Enters a Prompt
- Comparing Two Answers to One Question — Detecting Unstable Responses

Day 4: Controls, Thresholds and Accountability

- Consequence Tiering Matrix — Setting Checking Depth by Potential Harm
- Sampling Plans for High-Volume Output — Balancing Speed Against Assurance
- Human-in-the-Loop Sign-Off Rules — Naming Who Approves Each Output
- ISO/IEC 42001 and NIST AI RMF Controls — Anchoring Checks in Policy
- Fabrication Error Log — Feeding Lessons Back Into Prompts and Sources

Day 5: Practising the Answer Verification Protocol

- Live Draft Audit Exercise — Applying Claim Extraction Under Time Pressure
- Grounded Rewrite Challenge — Repairing a Flawed Briefing With Sources
- Tiering Your Own Assistant Uses — Mapping Department Outputs to Checks
- Peer Challenge on Sign-Off Rules — Testing Thresholds Against Real Cases
- Answer Verification Protocol Drafting — Final Version for Departmental Use

Applying Generative AI Hallucination Control After the Course

- Run a claim extraction audit on ten recent assistant-written documents and record each fabrication found in a shared error log.
- Agree consequence tiers and sign-off thresholds with the department head, then pilot the sampling plan on one recurring report.

Generative AI Hallucination Control Questions FAQ

What is generative AI hallucination control and why does it matter now?

It is the practice of catching and preventing false statements in AI-written output. It matters now because one invented figure in a client document can cost more trust than the assistant saved in time.

Which roles gain most from generative AI hallucination control training?

Professionals who draft with assistants and managers who approve those drafts gain most. Analysts, risk officers and source library curators also benefit.

How does a generative AI hallucination control course differ from a general AI tools course?

It focuses on proving outputs right rather than producing them faster. A general tools course teaches features for productivity; this one teaches claim checking, grounding and sign-off.

What do you bring back from generative AI hallucination control training?

You bring back a protocol your department can adopt at once. It sets which outputs need checking, how deep, and who signs off.

Generative AI Hallucination Control: The Way Forward

The Answer Verification Protocol first informs which assistant outputs may leave the department without a second reader. Drafters use it daily, managers use it at sign-off, and risk officers use its error log to judge whether controls work. After its first month, revisit the consequence tiers, because the outputs that seemed low risk often prove to carry the subtlest fabrications.