Faysal Aziz

From AI Fundamentals to Agentic Automation

A practical one-day introduction to building useful, controlled and secure AI workflows

FA-0002 · 1 day

Agentic & Generative AI · Automation, RPA & Power Platform

Artificial intelligence is moving beyond chat-based assistance into systems that can retrieve organisational knowledge, interpret incidents, modify software and coordinate business processes.

Introduction

Artificial intelligence is moving beyond chat-based assistance into systems that can retrieve organisational knowledge, interpret incidents, modify software and coordinate business processes. Understanding how these systems work is essential before allowing them to take actions inside a real environment.

This course begins with the foundations: what AI is, how machine learning and large language models work, how prompts are processed, and how Retrieval-Augmented Generation grounds responses in external information. It then introduces specification-driven development and the structure of an agentic automation harness capable of receiving an incident ticket, investigating the problem, proposing a fix, raising a pull request and progressing towards deployment under human approval.

A focused amount of theory is deliberately included because many AI risks—such as prompt injection, sensitive-data exposure, unreliable outputs, excessive permissions and untraceable automated decisions—are not immediately visible to users. Current guidance from NIST and OWASP emphasises that generative AI must be governed throughout its design, deployment and operation, rather than treated as an ordinary software feature.

Approximately the final third of the course will consist of instructor-led, hands-on demonstrations. Python will be used as the primary demonstration language, with Google Workspace presented as an alternative technology stack for building accessible business automations through services such as Gmail, Drive, Docs and Sheets.

The instructor brings more than 30 years of industry experience and will present practical, industry-demanded content rather than an academic treatment of the subject.

Learning outcomes

Participants will be able to:

- Explain AI, large language models, prompting and RAG
- Describe how AI models process and generate responses
- Understand specification-driven AI development
- Identify the main components of an agentic automation harness
- Map an incident-to-resolution AI workflow
- Recognise major AI security, privacy and operational risks
- Define basic tracking, logging and approval controls
- Understand how Python and Google Workspace can support AI automation

Prerequisites

- General understanding of IT, software or business workflows
- Basic awareness of incident tickets and source control is beneficial
- Programming experience is helpful but not required
- No previous AI experience is required

Training outline

1. Artificial Intelligence Fundamentals

- What AI is and how it works
- AI versus traditional software
- Machine learning and generative AI
- Large language models
- Tokens, prompts and context windows
- Model inference and probabilistic outputs
- Hallucinations and model limitations

2. Essential Generative AI Concepts

- Prompt engineering
- Embeddings and semantic search
- Retrieval-Augmented Generation

- Vector databases
- Tool and function calling
- AI workflows versus AI agents
- Model Context Protocol concepts

3. AI Specification-Driven Development

- Specifications as the source of truth
- Business requirements and system boundaries
- User stories and acceptance criteria
- Constraints and prohibited actions
- AI-generated code validation
- Specification versioning and traceability

4. Agentic AI Automation Harness

- Goals, plans, actions and observations
- Tool and API integration
- Workflow orchestration
- Memory and state management
- Human-in-the-loop approvals
- Failure handling and escalation
- Controlled execution boundaries

5. Incident-to-Resolution Workflow

- Automated ticket intake
- Incident classification
- Context and log retrieval
- Root-cause analysis support
- Proposed bug remediation
- Automated code modification
- Testing and validation
- Pull-request creation
- Human review and approval
- Controlled deployment
- Verification and rollback

- Ticket update and closure

6. AI Security, Risk and Governance

- Prompt injection
- Sensitive-data exposure
- Insecure AI outputs
- Excessive agent permissions
- Tool misuse and unauthorised actions
- Data privacy and retention
- Identity and access controls
- Human accountability
- Governance and approval boundaries

7. AI Tracking and Auditability

- Prompt and response logging
- Tool-call tracking
- Agent decision records
- Model and prompt version tracking
- Cost and token monitoring
- Accuracy and reliability evaluation
- Approval and override records
- Audit trails and operational dashboards

8. Instructor-Led Practical Demonstrations

Python AI workflow

- Model interaction
- Structured prompts
- Tool calling
- Ticket processing
- Approval-gated automation
- Logging and tracing

Agentic incident workflow

- Receive an incident ticket
- Analyse the issue

- Propose a fix
- Prepare a code change
- Raise a pull request
- Apply human approval
- Simulate deployment
- Record the audit trail

Google Workspace alternative stack

- Gmail ticket intake
- Google Sheets workflow tracking
- Google Drive knowledge retrieval
- Google Docs reporting
- Approval and notification workflows

Disclaimer

This training outline is provided as a general instructional guideline. The trainer reserves the right to amend, reorder, substitute or omit any topic where reasonably necessary to accommodate participant experience, technical availability, operational priorities or time constraints, without prior notice.

© 2026 Faysal Aziz · Kuala Lumpur