Faysal Aziz

Building Context-Aware AI Agents with Copilot Studio

Design smarter business agents by controlling context, knowledge, actions, and orchestration - 1 day

FA-0122 · 1 day

Agentic & Generative AI · Automation, RPA & Power Platform

Introduction

Once users are comfortable with generative AI and Microsoft Copilot, the next challenge is making AI useful beyond individual conversations. Enterprise agents need more than good prompts. They need the right context, access to appropriate knowledge, clearly defined instructions, controlled actions, and enough structure to behave consistently within a business process.

This one-day course moves beyond introductory AI concepts and concentrates on the practical design and construction of context-aware agents using Microsoft Copilot Studio. The emphasis is on understanding what information an agent receives, where that information comes from, how context is maintained during an interaction, and how knowledge, variables, instructions, tools, and business data influence the agent's decisions.

Current Copilot Studio capabilities increasingly center on generative orchestration, where agents can select among knowledge sources, tools, topics, and other capabilities according to the user's request rather than relying entirely on rigid conversational flows. Microsoft has also expanded its agent architecture around organizational knowledge, dynamic user context, connectors, tools, and Model Context Protocol integrations. This makes context design an important part of agent development: an effective agent must not simply have access to information; it must receive the right information at the right point and operate within appropriate boundaries.

The course therefore focuses heavily on **context engineering and agent construction**, rather than introductory prompting or general AI theory. Participants will work at the level of agent purpose, instructions, knowledge sources, variables, tools, orchestration, user context, and testing. Copilot Studio's

current knowledge model can incorporate sources such as SharePoint, Dataverse, Azure AI Search, enterprise connectors, and other organizational information, while orchestration determines how those resources are selected and used.

The instructor has **over 30 years of industry experience** and will approach the subject from the perspective of building workable business solutions rather than delivering an academic treatment of AI. Because this is a one-day intermediate course, coverage is intentionally concentrated on the agent capabilities that participants can realistically understand and apply within the available time.

Learning outcomes

By the end of the course, participants should be able to:

- Design the purpose, scope, and boundaries of a business AI agent.
- Explain the difference between instructions, conversational context, knowledge, and tools.
- Structure agent context to improve relevance and consistency.
- Configure appropriate knowledge sources for an agent.
- Use variables and runtime information to influence agent behavior.
- Understand generative orchestration and agent decision-making.
- Configure tools and actions for business operations.
- Understand how connectors and external capabilities extend an agent.
- Design appropriate human intervention and control points.
- Test and troubleshoot agent behavior using Copilot Studio capabilities.
- Recognize key security and governance considerations when deploying business agents.

Prerequisites

- Good understanding of generative AI and LLMs.
- Working experience using Microsoft Copilot or a comparable generative AI platform.
- General understanding of business processes and workflow automation.
- Working experience with Microsoft Power Platform.
- No advanced programming experience is required.

Training outline

1. Designing Business AI Agents

Agent Purpose and Scope

1. Business objective
2. User requirements
3. Agent responsibilities
4. Operational boundaries
5. Human responsibility

Agent Architecture

1. Instructions
2. Knowledge
3. Context
4. Tools
5. Topics and triggers
6. Orchestration

2. Context Engineering for AI Agents

Understanding Agent Context

1. User context
2. Conversation context
3. Business context
4. Runtime context
5. Organizational context

Managing Context

1. Variables
2. Input and output values
3. Conversation state
4. Context persistence

5. Dynamic context

Context and Agent Instructions

1. Agent-level instructions
2. Behavioral boundaries
3. Knowledge references
4. Tool references
5. Context-aware decision making

3. Knowledge and Grounding

Knowledge Architecture

1. Agent knowledge sources
2. Organizational knowledge
3. SharePoint
4. Dataverse
5. Enterprise data sources

Knowledge Source Design

1. Source selection
2. Knowledge descriptions
3. Information boundaries
4. User permissions
5. Grounded responses

Dynamic Knowledge and Context

1. Context-sensitive retrieval
2. Runtime knowledge selection
3. User-specific information
4. Knowledge source variables
5. Organizational context

4. Building Agents in Microsoft Copilot Studio

Agent Configuration

1. Agent creation
2. Instructions
3. Knowledge configuration
4. Conversation behavior
5. Generative AI settings

Generative Orchestration

1. Orchestration concepts
2. Knowledge selection
3. Tool selection
4. Topic selection
5. Multi-step reasoning

Topics and Structured Behavior

1. Topic design
2. Triggers
3. Variables
4. Conditions
5. Agent handoffs

5. Tools, Actions and Business Integration

Agent Tools

1. Tool concepts
2. Input parameters
3. Output values
4. Tool descriptions
5. Tool selection

Extending Agent Capabilities

1. Power Automate flows
2. Connectors
3. APIs
4. Model Context Protocol
5. External services

Action Governance

1. Authentication
2. Authorization
3. User confirmation
4. Transaction boundaries
5. Error handling

6. Testing, Control and Deployment

Agent Testing

1. Conversation testing
2. Activity maps
3. Knowledge testing
4. Tool execution
5. Context validation

Agent Quality and Control

1. Response grounding
2. Instruction conflicts
3. Context failures
4. Tool-selection failures
5. Escalation and human intervention

Security and Governance

1. Data access

2. Connector controls
3. Environment governance
4. Publishing controls
5. Monitoring and lifecycle management

Disclaimer

This course outline is provided as a proposed instructional framework and is intended to guide delivery within the stated one-day duration. The trainer reserves the right to modify, reorder, consolidate, expand, reduce, or substitute topics where reasonably necessary in response to participant experience, available time, technological changes, platform availability, or instructional priorities, without prior notice.

© 2026 Faysal Aziz · Kuala Lumpur