

01	A	B	C	D	E
02	A	B	C	D	E
03	A	B	C	D	E
04	A	B	C	D	E
05	A	B	C	D	E
06	A	B	C	D	E
07	A	B	C	D	E
08	A	B	C	D	E
09	A	B	C	D	E
10	A	B	C	D	E
11	A	B	C	D	E
12	A	B	C	D	E
13	A	B	C	D	E

14	A	B	C	D	E
15	A	B	C	D	E
16	A	B	C	D	E
17	A	B	C	D	E
18	A	B	C	D	E
19	A	B	C	D	E
20	A	B	C	D	E
21	A	B	C	D	E
22	A	B	C	D	E
23	A	B	C	D	E
24	A	B	C	D	E
25	A	B	C	D	E
26	A	B	C	D	E

27	A	B	C	D	E
28	A	B	C	D	E
29	A	B	C	D	E
30	A	B	C	D	E
31	A	B	C	D	E
32	A	B	C	D	E
33	A	B	C	D	E
34	A	B	C	D	E
35	A	B	C	D	E
36	A	B	C	D	E
37	A	B	C	D	E
38	A	B	C	D	E
39	A	B	C	D	E

PRACTICE QUESTIONS

OutSystems

AISOJ

OutSystems Agentic AI Specialization (JA)

EDITION

Demo

QUESTIONS

9

FORMAT

Limited Edition PDF

Before you begin

What this is

9 practice questions for OutSystems AIS0J , each with its answer key and an explanation. Answer a question before you read its key — the explanation is worth more once you have committed.

If something looks wrong

Send us three things and we will check it:

- the exam code, AIS0J
- the question number
- the email address on your order

Write to **support@mometrix.net**. We reply within one business day.

QUESTION 1

MULTIPLE CHOICE

When designing an AI Agent in OutSystems that needs to retrieve information from a knowledge base before responding, which components should be included in the agent workflow? (Choose two.)

- ☒ **A** An AI Agent Server Action that defines the agent's instructions and behavior
- ☒ **B** A Knowledge Base entity pre-populated with source content
- ☐ **C** A traditional OutSystems Timer scheduled to run every minute
- ☐ **D** An Email Server module to dispatch responses to end users

ANSWER A - B

WHY An AI Agent in OutSystems is implemented as a Server Action that defines the agent's instructions, model, and behavior (A). To ground the agent with domain content, the agent references a Knowledge Base that has been populated with source documents or data (B). A Timer (C) and an Email module (D) are unrelated to the agent's retrieval-augmented generation pattern and are not required components.

QUESTION 2

MULTIPLE CHOICE

Which of the following are valid considerations when choosing an LLM for an OutSystems AI Agent? (Choose two.)

- ☒ **A** Cost per token and rate limits of the model provider
- ☒ **B** Whether the model supports the languages and context length required by the use case
- ☐ **C** The screen resolution of the end user's device
- ☐ **D** Whether the model can be embedded in a local OutSystems mobile package

ANSWER A - B

WHY When selecting an LLM, the agent designer must weigh cost and rate limits of the provider (A) and confirm the model supports the required languages and context window for the use case (B). Screen resolution (C) is irrelevant to LLM selection, and LLMs are not packaged inside an OutSystems mobile binary (D); they are invoked via service calls to model providers.

QUESTION 3

SINGLE CHOICE

In an OutSystems AI Agent flow, what is the primary role of the system prompt?

- ☐ A To define the model provider's API endpoint URL
- ☒ B To specify the agent's persona, goals, constraints, and how it should use its tools
- ☐ C To encrypt the conversation history stored in the database
- ☐ D To allocate GPU resources for the inference engine

ANSWER B

WHY The system prompt establishes the agent's identity, objectives, behavioral guidelines, and instructions for invoking its tools (B). It does not configure API endpoints (A), handle encryption (C), or manage infrastructure resources (D).

QUESTION 4

SINGLE CHOICE

Which prompt-engineering technique is most appropriate when an OutSystems AI Agent must consistently produce output in a strict JSON schema for downstream processing?

- ☐ A Free-form conversation prompting
- ☒ B Structured output prompting with explicit schema and example payload
- ☐ C Prompt injection with adversarial examples
- ☐ D Few-shot image captioning prompting

ANSWER B

WHY When deterministic, machine-readable output is required, structured output prompting that explicitly defines the JSON schema and provides an example payload is the correct technique (B). Free-form prompts (A) cannot guarantee schema adherence, prompt injection (C) is a security concern rather than a technique to enforce structure, and image captioning (D) is unrelated to JSON output.

QUESTION 5

SINGLE CHOICE

An OutSystems AI Agent needs up-to-date product catalog data that lives in an external REST API. What is the recommended OutSystems approach to make that data available to the agent at runtime?

- ☐ A Embed the catalog as a static text block inside the system prompt
- ☒ B Expose the data through a Server Action or REST endpoint that the agent can invoke as a tool
- ☐ C Hardcode the catalog values into the OutSystems client screen JavaScript
- ☐ D Schedule a BPT process to email the catalog to the LLM provider daily

ANSWER B

WHY OutSystems AI Agents can call external logic through tools. Wrapping the REST call in a Server Action or REST endpoint and registering it as a tool lets the agent fetch fresh data on demand (B). Static prompt content (A) goes stale, client-side hardcoding (C) bypasses the agent, and emailing data (D) is not a supported integration pattern.

QUESTION 6

SINGLE CHOICE

Which practice best reduces the risk of prompt injection when an OutSystems AI Agent ingests untrusted user input?

- ☐ A Disabling HTTPS on the agent endpoint to simplify debugging
- ☒ B Separating untrusted content from instructions and validating/quoting it before it reaches the model
- ☐ C Increasing the model's temperature to zero to force deterministic replies
- ☐ D Sharing the system prompt publicly so users know how the agent works

ANSWER B

WHY Treating untrusted input as data rather than instructions — by clearly separating it from system/user instructions and validating or delimiting it — is the recommended mitigation against prompt injection (B). Disabling HTTPS (A) weakens security, setting temperature to zero (C) affects randomness rather than injection safety, and publishing the system prompt (D) makes injection easier.

QUESTION 7

SINGLE CHOICE

When evaluating an OutSystems AI Agent before deployment, which approach most directly measures whether the agent chooses the correct tool for a given user request?

- ☐ A Visual inspection of the agent's screen layout
- ☒ B Running a curated test set of prompts and asserting on the tools invoked and final answers
- ☐ C Manually clicking through the agent in Production with real users
- ☐ D Reviewing the OutSystems Service Center license usage dashboard

ANSWER B

WHY A representative evaluation set of prompts with assertions on the tools selected and the answers produced is the standard way to measure tool-selection accuracy (B). Visual layout (A) does not test reasoning, testing in Production with real users (C) is risky and unstructured, and the license dashboard (D) reports consumption, not quality.

QUESTION 8

SINGLE CHOICE

A business stakeholder asks whether a use case is a good fit for an OutSystems AI Agent. Which characteristic most strongly indicates that an AI Agent is appropriate?

- ☐ A The task is fully deterministic and can be expressed as a fixed decision tree
- ☒ B The task requires understanding natural language, reasoning over context, and orchestrating multiple steps or tools
- ☐ C The task only involves rendering a static report from a database
- ☐ D The task requires no user interaction at all

ANSWER B

WHY AI Agents shine when the task involves natural language understanding, contextual reasoning, and multi-step orchestration across tools (B). Deterministic rules (A) are better handled by traditional logic, static reporting (C) needs no LLM, and zero-interaction tasks (D) do not benefit from conversational reasoning.

QUESTION 9

SINGLE CHOICE

Which statement best describes 'few-shot prompting' as used with an OutSystems AI Agent?

- ☒ **A** Providing several example inputs and desired outputs within the prompt to steer the model's behavior
- ☐ **B** Running the model multiple times with different random seeds in parallel
- ☐ **C** Training the underlying model on a small custom dataset
- ☐ **D** Limiting the conversation history to only the most recent few user turns

ANSWER A

WHY Few-shot prompting supplies the model with a handful of in-context examples that demonstrate the desired input/output pattern (A). It is not parallel inference (B), model fine-tuning (C), or a truncation strategy (D).