

What must the AI do?

Create something new

Pursue a goal and take actions

Understand written language

Support assistant

Hear or speak

Interpret visual content

Extract structured information

Key message

The workload is the kind of problem being solved.

The six workload families

Workload	Defining capability
Generative AI	Creates new content
Agentic AI	Pursues goals through decisions and actions
Text analysis	Discovers meaning in language
Speech	Processes or produces spoken language
Computer vision	Interprets or creates visual content
Information extraction	Converts unstructured content into structured data

Ask: “What transformation does the application perform?”

Generative AI

Produces new content based on instructions and context.

Common outputs

- Answers and summaries
- Emails and reports
- Images
- Code
- Translations or rewritten text

Scenario signal

“Generate,” “compose,” “draft,” “rewrite,” or “create”

Example

A support assistant drafts a personalized answer using the customer’s question and company documentation.

Exam trap

Summarization can appear under both generative AI and text analysis. Use the scenario’s emphasis:

Producing a fluent new response → Generative AI

Identifying or analyzing meaning in existing text → Text analysis


Agentic AI

Uses a model to decide what to do next and interact with tools or systems.

Possible actions

- Search a knowledge source
- Call an API
- Update a record
- Create a support ticket
- Ask for missing information
- Hand work to a human

Typical loop



Example

The assistant checks an order, requests a replacement, updates the CRM, and tells the customer what happened.

Exam distinction

Produces an answer only → Generative AI

Chooses and performs actions toward a goal → Agentic AI

Text analysis

Finds meaning, patterns, or classifications in text.

Core techniques

- Keyword extraction
- Entity detection
- Sentiment analysis
- Language detection
- Classification
- Summarization

Scenario signal

The desired result describes or labels existing text.

Example

Analyze a customer message to determine:

Negative sentiment · Product name · Account number · Cancellation intent

Speech

Processes or produces spoken language.

Capabilities

- Speech recognition: speech → text
- Speech synthesis: text → speech
- Speech translation: speech → another language
- Spoken interaction with multimodal models

Memory hook

If audio represents human language, think speech.

Example

A caller describes a problem aloud, and the assistant responds using a synthesized voice.

Exam trap

Analyzing nonverbal audio—such as alarms, machinery, or environmental sounds—may be information extraction rather than speech.

Computer vision

Interprets images or video—or generates new visual outputs.

Common capabilities

- Image classification
- Object detection
- Image description
- Visual question answering
- Optical character recognition
- Image generation

Scenario signal

The answer depends on what is visually present.

Example

A customer uploads a photograph, and the assistant identifies visible damage to the product.

Exam trap

Determine what an image depicts → Computer vision

Pull named fields from the image → Information extraction

Information extraction

Converts unstructured or semi-structured content into structured fields.

Supported sources can include

- Documents and forms
- Images
- Audio
- Video
- Text

Example output

Customer: Maria Chen
Invoice: 10482
Date: 2026-09-12
Total: \$847.50

Example output

Structured fields captured from an invoice.

Key distinction

Explain what a document says → Text analysis or generative AI

Capture specific fields from it → Information extraction

The customer-support assistant

1 Transcribe the caller

Speech

2 Detect frustration

Text analysis

3 Inspect a damage photo

Computer vision

4 Capture invoice fields

Information extraction

5 Draft the response

Generative AI

6 Issue a refund through an API

Agentic AI

Key takeaway

Workloads are composable capabilities—not mutually exclusive application categories.

Classify the transformation

1. Find the verb
2. Identify the input
3. Inspect the expected output

Scenario language	Likely workload
Draft an answer	Generative AI
Complete the task	Agentic AI
Detect sentiment	Text analysis
Transcribe a call	Speech
Identify an object	Computer vision
Populate invoice fields	Information extraction

Do not classify the industry. Classify what the AI must do.

A support system must read emailed invoices and store the invoice number, supplier, date, and total in a database.

Which workload is primary?

A. Generative AI

B. Text analysis

C. Computer vision

D. Information extraction

Answer reveal

D. Information extraction

The decisive requirement is converting document content into predefined structured fields. The presence of text does not automatically make the scenario text analysis.