

1. Count the Tokens

You type the command “fetch the latest price data” into a task-planning agent. The agent splits the text on spaces, counting each word as one token. How many tokens does it count?

Step 1: Write each token in a box

fetch	the	latest	price	data			
-------	-----	--------	-------	------	--	--	--

Step 2: Count the filled boxes

5 tokens

Practice 1

You send the sentence “Agents think step by step” to an AI agent. The agent splits the text on spaces, counting each word as one token. How many tokens does it count?

Step 1: Write each token in a box

--	--	--	--	--	--	--	--

Step 2: Count the filled boxes

--

Practice 2

You send the instruction “plan and run the next step” to an automation agent. The agent splits the text on spaces, counting each word as one token. How many tokens does it count?

Step 1: Write each token in a box

--	--	--	--	--	--	--	--

Step 2: Count the filled boxes

--

2. Subword Splitting

A tokenizer splits “preprocessing” into three pieces: “pre”, “process”, and “ing”. The agent receives the pipeline command “run preprocessing now”. How many tokens does it count?

Step 1: Write each token in a box

run	pre	process	ing	now			
-----	-----	---------	-----	-----	--	--	--

Step 2: Count the filled boxes

5 tokens

Practice 1

A tokenizer splits “embeddings” into “embed” and “dings”, and “retrieval” into “retriev” and “al”. The agent receives the query “run embeddings retrieval”. How many tokens does it count?

Step 1: Write each token in a box

--	--	--	--	--	--	--	--

Step 2: Count the filled boxes

--

Practice 2

A tokenizer splits “debugging” into “debug” and “ging”, and “codebase” into “code” and “base”. The agent receives the task “start debugging codebase issues”. How many tokens does it count?

Step 1: Write each token in a box

--	--	--	--	--	--	--	--

Step 2: Count the filled boxes

--

3. Punctuation Counts

An AI agent counts each punctuation mark as its own token. The agent receives “Hello, world!”. How many tokens does it count?

Step 1: Write each token in a box

Hello	,	world	!		
-------	---	-------	---	--	--

Step 2: Count the filled boxes

4 tokens

Practice 1

An AI agent counts each punctuation mark as its own token. The agent receives “Wait, what?”. How many tokens does it count?

Step 1: Write each token in a box

--	--	--	--	--	--

Step 2: Count the filled boxes

--

Practice 2

An AI agent counts each punctuation mark as its own token. The agent receives “Stop, look, and listen!”. How many tokens does it count?

Step 1: Write each token in a box

Step 2: Count the filled boxes

--

4. Tokens per Word

An AI agent processes a 20-word passage of English text and counts 26 tokens. What is the token-to-word ratio?

Step 1: Divide tokens by words

$$26/20 = 1.3 \text{ tokens per word}$$

Practice 1

An AI agent processes a 24-word passage of Spanish text and counts 36 tokens. What is the token-to-word ratio?

Step 1: Divide tokens by words

Practice 2

An AI agent processes a 15-word passage of Hindi text and counts 45 tokens. What is the token-to-word ratio?

Step 1: Divide tokens by words

5. Will It Fit?

An agent has a 4,000-token context window. The system prompt and conversation history already use 1,200 tokens. The document to read is 55 paragraphs, each about 50 tokens. Can the agent read the document in this call?

Step 1: Space available for the document

$$4,000 - 1,200 = 2,800 \text{ tokens}$$

Step 2: Estimate the document

$$55 \times 50 = 2,750 \text{ tokens}$$

Step 3: Does it fit?

$$2,750 < 2,800 \text{ — yes, but barely.}$$

Practice 1

An agent has a 4,000-token context window. The system prompt and conversation history already use 600 tokens. The document to read is 8 pages, each about 500 tokens. Can the agent read the document in this call?

Step 1: Space available for the document

Step 2: Estimate the document

Step 3: Does it fit?

5. Will It Fit? (continued)

Practice 2

An agent has a 4,000-token context window. The system prompt and conversation history already use 1,000 tokens. The document to read is 18 slides, each about 150 tokens. Can the agent read the document in this call?

Step 1: Space available for the document

Step 2: Estimate the document

Step 3: Does it fit?