

# DEVELOPING WITH AI APIS

Module 7 — Harwell Prompt Engineering

# LEARNING OBJECTIVES

By the end of this module you will be able to:



# LEARNING OBJECTIVES

By the end of this module you will be able to:

- Explain the “delta” between standard REST APIs and LLM APIs: stateful vs. stateless, non-determinism, streaming

- 

- 

- 

-

# LEARNING OBJECTIVES

By the end of this module you will be able to:

- Explain the “delta” between standard REST APIs and LLM APIs: stateful vs. stateless, non-determinism, streaming
- Describe why LLM APIs differ from typical REST: stateful conversations, non-deterministic responses

- 

- 

-

# LEARNING OBJECTIVES

By the end of this module you will be able to:

- Explain the “delta” between standard REST APIs and LLM APIs: stateful vs. stateless, non-determinism, streaming
- Describe why LLM APIs differ from typical REST: stateful conversations, non-deterministic responses
- Control non-determinism: temperature and top-p; when to use low vs. high
- 
-

# LEARNING OBJECTIVES

By the end of this module you will be able to:

- Explain the “delta” between standard REST APIs and LLM APIs: stateful vs. stateless, non-determinism, streaming
- Describe why LLM APIs differ from typical REST: stateful conversations, non-deterministic responses
- Control non-determinism: temperature and top-p; when to use low vs. high
- Handle streaming responses and implications for UX (tokens/sec, buffering)
-

# LEARNING OBJECTIVES

By the end of this module you will be able to:

- Explain the “delta” between standard REST APIs and LLM APIs: stateful vs. stateless, non-determinism, streaming
- Describe why LLM APIs differ from typical REST: stateful conversations, non-deterministic responses
- Control non-determinism: temperature and top-p; when to use low vs. high
- Handle streaming responses and implications for UX (tokens/sec, buffering)
- Understand cost drivers: token economics (input vs. output), and how to reason about cost when designing features

# BRIDGE FROM MODULE 6

What we learned last time:

- 
-

# BRIDGE FROM MODULE 6

What we learned last time:

- MCP connects AI to live systems
-

# BRIDGE FROM MODULE 6

What we learned last time:

- MCP connects AI to live systems
- MCP uses APIs under the hood

## Frame explicitly:

- 
- 
- 

**Today:** Understand how LLM APIs work differently from standard REST.

## Frame explicitly:

- This is the **delta** — what's different from standard REST APIs
- 
- 

**Today:** Understand how LLM APIs work differently from standard REST.

## Frame explicitly:

- This is the **delta** — what's different from standard REST APIs
- Not a full SDK deep dive
- 

**Today:** Understand how LLM APIs work differently from standard REST.

## Frame explicitly:

- This is the **delta** — what's different from standard REST APIs
- Not a full SDK deep dive
- Focus on understanding differences

**Today:** Understand how LLM APIs work differently from standard REST.

# THE PROBLEM: LLM APIS ARE DIFFERENT

Standard REST expectations:

- 
- 
- 
-

# THE PROBLEM: LLM APIS ARE DIFFERENT

Standard REST expectations:

- Stateless requests
- 
- 
-

# THE PROBLEM: LLM APIS ARE DIFFERENT

Standard REST expectations:

- Stateless requests
- Deterministic responses
- 
-

# THE PROBLEM: LLM APIS ARE DIFFERENT

**Standard REST expectations:**

- Stateless requests
- Deterministic responses
- Simple request/response
-

# THE PROBLEM: LLM APIS ARE DIFFERENT

**Standard REST expectations:**

- Stateless requests
- Deterministic responses
- Simple request/response
- Predictable costs

## LLM APIs break these:

- 
- 
- 
-

## LLM APIs break these:

- **X** Stateful conversations vs. stateless
- 
- 
-

## LLM APIs break these:

- **X** Stateful conversations vs. stateless
- **X** Non-deterministic responses
- 
-

## LLM APIs break these:

- **X** Stateful conversations vs. stateless
- **X** Non-deterministic responses
- **X** Streaming chunks vs. complete response
-

## LLM APIs break these:

- **X** Stateful conversations vs. stateless
- **X** Non-deterministic responses
- **X** Streaming chunks vs. complete response
- **X** Token-based costs vs. request-based

## Why it matters:

- 
- 
- 
-

## Why it matters:

- **X** Expect stateless → conversations won't work
- 
- 
-

## Why it matters:

- **X** Expect stateless → conversations won't work
- **X** Expect deterministic → confused by variation
- 
-

## Why it matters:

- **X** Expect stateless → conversations won't work
- **X** Expect deterministic → confused by variation
- **X** Don't handle streaming → poor UX
-

## Why it matters:

- **X** Expect stateless → conversations won't work
- **X** Expect deterministic → confused by variation
- **X** Don't handle streaming → poor UX
- **X** Ignore cost → high bills

# STATEFUL VS. STATELESS: CONVERSATIONS

Standard REST:

- 
- 
- 
-

# STATEFUL VS. STATELESS: CONVERSATIONS

Standard REST:

- **Stateless:** Each request is independent
- 
- 
-

# STATEFUL VS. STATELESS: CONVERSATIONS

## Standard REST:

- **Stateless:** Each request is independent
- Example: GET /users/123 → returns user data
- 
-

# STATEFUL VS. STATELESS: CONVERSATIONS

## Standard REST:

- **Stateless:** Each request is independent
- Example: GET /users/123 → returns user data
- No memory between requests
-

# STATEFUL VS. STATELESS: CONVERSATIONS

## Standard REST:

- **Stateless:** Each request is independent
- Example: GET /users/123 → returns user data
- No memory between requests
- Simple, predictable

## LLM APIs:

- 
- 
- Request 1: “What is Spring Boot?”
- Request 2: “How do I use it?” (refers to previous)
- 
-

## LLM APIs:

- **Stateful:** Conversations maintain context
- - Request 1: “What is Spring Boot?”
  - Request 2: “How do I use it?” (refers to previous)
- 
-

## LLM APIs:

- **Stateful:** Conversations maintain context
- Example:
  - Request 1: “What is Spring Boot?”
  - Request 2: “How do I use it?” (refers to previous)
- 
-

## LLM APIs:

- **Stateful:** Conversations maintain context
- Example:
  - Request 1: “What is Spring Boot?”
  - Request 2: “How do I use it?” (refers to previous)
- **Conversation ID or message list:** Maintains context
-

## LLM APIs:

- **Stateful:** Conversations maintain context
- Example:
  - Request 1: “What is Spring Boot?”
  - Request 2: “How do I use it?” (refers to previous)
- **Conversation ID or message list:** Maintains context
- More complex, but enables natural conversations

# HOW STATEFUL WORKS

## Option 1: Conversation ID

- 
- 
- 
-

# HOW STATEFUL WORKS

## Option 1: Conversation ID

- First request returns conversation\_id
- 
- 
-

# HOW STATEFUL WORKS

## Option 1: Conversation ID

- First request returns conversation\_id
- Subsequent requests include conversation\_id
- 
-

# HOW STATEFUL WORKS

## Option 1: Conversation ID

- First request returns conversation\_id
- Subsequent requests include conversation\_id
- API maintains conversation state
-

# HOW STATEFUL WORKS

## Option 1: Conversation ID

- First request returns conversation\_id
- Subsequent requests include conversation\_id
- API maintains conversation state
-  Simpler for client

## Option 2: Message list

- 
- 
- 
-

## Option 2: Message list

- You maintain message history
- 
- 
-

## Option 2: Message list

- You maintain message history
- Send full message list with each request
- 
-

## Option 2: Message list

- You maintain message history
- Send full message list with each request
- You control state
-

## Option 2: Message list

- You maintain message history
- Send full message list with each request
- You control state
-  More control

## When to use:



## When to use:

- **Stateless (one-off):** Simple Q&A, no context needed
-

## When to use:

- **Stateless (one-off):** Simple Q&A, no context needed
- **Stateful (conversation):** Multi-turn dialogue, follow-up questions

# NON-DETERMINISM: THE PROBLEM

Same prompt → different answers

- 
-

# NON-DETERMINISM: THE PROBLEM

Same prompt → different answers

- “Why is this happening?”
-

# NON-DETERMINISM: THE PROBLEM

Same prompt → different answers

- “Why is this happening?”
- “How do I control it?”

## Why non-deterministic?

- 
- 
- 
-

## Why non-deterministic?

- LLMs are probabilistic, not deterministic
- 
- 
-

## Why non-deterministic?

- LLMs are probabilistic, not deterministic
- They sample from probability distributions
- 
-

## Why non-deterministic?

- LLMs are probabilistic, not deterministic
- They sample from probability distributions
- Same input can produce different outputs
-

## Why non-deterministic?

- LLMs are probabilistic, not deterministic
- They sample from probability distributions
- Same input can produce different outputs
- This is by design (creativity, variety)

# CONTROLLING NON- DETERMINISM

Temperature: Controls randomness

- 
- 
-

# CONTROLLING NON- DETERMINISM

Temperature: Controls randomness

- **Low (0-0.3):** More deterministic, focused
- 
-

# CONTROLLING NON- DETERMINISM

Temperature: Controls randomness

- **Low (0-0.3):** More deterministic, focused
- **Medium (0.5-0.7):** Balanced
-

# CONTROLLING NON- DETERMINISM

Temperature: Controls randomness

- **Low (0-0.3):** More deterministic, focused
- **Medium (0.5-0.7):** Balanced
- **High (0.8-1.0):** More creative, varied

## Top-P (nucleus sampling): Controls diversity

- 
-

**Top-P (nucleus sampling):** Controls diversity

- **Low (0.1-0.3):** Focused on most likely tokens
-

## **Top-P (nucleus sampling): Controls diversity**

- **Low (0.1-0.3):** Focused on most likely tokens
- **High (0.9-1.0):** More diverse options

## When you need reproducibility:

- 
- 
-

## When you need reproducibility:

- Set temperature=0
- 
-

## When you need reproducibility:

- Set temperature=0
- Use same seed
-

## When you need reproducibility:

- Set temperature=0
- Use same seed
- More deterministic (not perfect)

# WHEN TO USE WHAT TEMPERATURE

Low temperature (0-0.3):

- 
- 
- 
-

# WHEN TO USE WHAT TEMPERATURE

Low temperature (0-0.3):

-  Code generation
- 
- 
-

# WHEN TO USE WHAT TEMPERATURE

Low temperature (0-0.3):

-  Code generation
-  Factual answers
- 
-

# WHEN TO USE WHAT TEMPERATURE

Low temperature (0-0.3):

-  Code generation
-  Factual answers
-  Tests
-

# WHEN TO USE WHAT TEMPERATURE

Low temperature (0-0.3):

-  Code generation
-  Factual answers
-  Tests
-  When you need consistency

## Medium temperature (0.5-0.7):

- 
- 

## High temperature (0.8-1.0):

- 
- 
-

## Medium temperature (0.5-0.7):

-  General use
- 

## High temperature (0.8-1.0):

- 
- 
-

## Medium temperature (0.5-0.7):

-  General use
-  Balanced responses

## High temperature (0.8-1.0):

- 
- 
-

## Medium temperature (0.5-0.7):

-  General use
-  Balanced responses

## High temperature (0.8-1.0):

-  Creative writing
- 
-

## Medium temperature (0.5-0.7):

-  General use
-  Balanced responses

## High temperature (0.8-1.0):

-  Creative writing
-  Brainstorming
-

## Medium temperature (0.5-0.7):

-  General use
-  Balanced responses

## High temperature (0.8-1.0):

-  Creative writing
-  Brainstorming
-  Varied responses

# STREAMING: CHUNKED RESPONSES

The problem:

- 
- 
- 
-

# STREAMING: CHUNKED RESPONSES

The problem:

- Without streaming: Wait for full response
- 
- 
-

# STREAMING: CHUNKED RESPONSES

The problem:

- Without streaming: Wait for full response
- **X** Slow perceived performance
- 
-

# STREAMING: CHUNKED RESPONSES

The problem:

- Without streaming: Wait for full response
- **X** Slow perceived performance
- **X** User sees nothing until complete
-

# STREAMING: CHUNKED RESPONSES

The problem:

- Without streaming: Wait for full response
- **X** Slow perceived performance
- **X** User sees nothing until complete
- **X** Poor UX for long responses

## The solution:

- 
- 
- 
-

## The solution:

- **Streaming:** Responses come in chunks (tokens)
- 
- 
-

## The solution:

- **Streaming:** Responses come in chunks (tokens)
-  User sees progress immediately
- 
-

## The solution:

- **Streaming:** Responses come in chunks (tokens)
-  User sees progress immediately
-  Better perceived performance
-

## The solution:

- **Streaming:** Responses come in chunks (tokens)
-  User sees progress immediately
-  Better perceived performance
-  Can cancel if needed

# HOW STREAMING WORKS

Server-Sent Events (SSE) or WebSocket

- 
- 
- 
-

# HOW STREAMING WORKS

Server-Sent Events (SSE) or WebSocket

- Tokens arrive incrementally
- 
- 
-

# HOW STREAMING WORKS

## Server-Sent Events (SSE) or WebSocket

- Tokens arrive incrementally
- Client buffers and displays as received
- 
-

# HOW STREAMING WORKS

## Server-Sent Events (SSE) or WebSocket

- Tokens arrive incrementally
- Client buffers and displays as received
- **Tokens/sec**: Measure of streaming speed
-

# HOW STREAMING WORKS

## Server-Sent Events (SSE) or WebSocket

- Tokens arrive incrementally
- Client buffers and displays as received
- **Tokens/sec**: Measure of streaming speed
- **Buffering**: Client may buffer before displaying

## Implementation considerations:

- 
- 
- 
-

## Implementation considerations:

- Handle partial responses
- 
- 
-

## Implementation considerations:

- Handle partial responses
- Display incrementally
- 
-

## Implementation considerations:

- Handle partial responses
- Display incrementally
- Handle errors mid-stream
-

## Implementation considerations:

- Handle partial responses
- Display incrementally
- Handle errors mid-stream
- Buffer for smooth display

# COST: TOKEN ECONOMICS

The problem:

- 
- 
- 
-

# COST: TOKEN ECONOMICS

The problem:

- Standard APIs: Cost per request (simple)
- 
- 
-

# COST: TOKEN ECONOMICS

The problem:

- Standard APIs: Cost per request (simple)
- LLM APIs: Cost per token (complex)
- 
-

# COST: TOKEN ECONOMICS

The problem:

- Standard APIs: Cost per request (simple)
- LLM APIs: Cost per token (complex)
- **X** Easy to underestimate costs
-

# COST: TOKEN ECONOMICS

The problem:

- Standard APIs: Cost per request (simple)
- LLM APIs: Cost per token (complex)
- **X** Easy to underestimate costs
- **X** Costs scale with usage

## Token economics:

- 
- 
- 
-

## Token economics:

- **Input tokens:** What you send (prompt + context)
- 
- 
-

## Token economics:

- **Input tokens:** What you send (prompt + context)
- **Output tokens:** What you receive (response)
- 
-

## Token economics:

- **Input tokens:** What you send (prompt + context)
- **Output tokens:** What you receive (response)
- **Pricing:** Usually per 1K tokens
-

## Token economics:

- **Input tokens:** What you send (prompt + context)
- **Output tokens:** What you receive (response)
- **Pricing:** Usually per 1K tokens
- **Input vs. output:** Output often more expensive

# **COST DRIVERS**

**What affects cost:**

- 
- 
- 
-

# COST DRIVERS

What affects cost:

- **Length:** Longer prompts/responses = more tokens
- 
- 
-

# COST DRIVERS

What affects cost:

- **Length:** Longer prompts/responses = more tokens
- **Frequency:** More requests = higher cost
- 
-

# COST DRIVERS

What affects cost:

- **Length:** Longer prompts/responses = more tokens
- **Frequency:** More requests = higher cost
- **Model:** Different models have different prices
-

# COST DRIVERS

What affects cost:

- **Length:** Longer prompts/responses = more tokens
- **Frequency:** More requests = higher cost
- **Model:** Different models have different prices
- **Context:** Including context increases input tokens

## Cost optimization:

- 
- 
- 
- 
-

## Cost optimization:

-  Shorter prompts (remove unnecessary context)
- 
- 
- 
-

## Cost optimization:

-  Shorter prompts (remove unnecessary context)
-  Cache responses when possible
- 
- 
-

## Cost optimization:

-  Shorter prompts (remove unnecessary context)
-  Cache responses when possible
-  Use cheaper models when appropriate
- 
-

## Cost optimization:

-  Shorter prompts (remove unnecessary context)
-  Cache responses when possible
-  Use cheaper models when appropriate
-  Monitor token usage
-

## Cost optimization:

-  Shorter prompts (remove unnecessary context)
-  Cache responses when possible
-  Use cheaper models when appropriate
-  Monitor token usage
-  Set usage limits

# ROUGH COST ESTIMATES

Examples:

- 
- 
- 
-

# ROUGH COST ESTIMATES

## Examples:

- 1000 tokens  $\approx$  750 words
- 
- 
-

# ROUGH COST ESTIMATES

## Examples:

- 1000 tokens  $\approx$  750 words
- Example pricing: \$0.01 per 1K input, \$0.03 per 1K output
- 
-

# ROUGH COST ESTIMATES

## Examples:

- 1000 tokens  $\approx$  750 words
- Example pricing: \$0.01 per 1K input, \$0.03 per 1K output
- Typical request: 500 input + 500 output = ~\$0.02
-

# ROUGH COST ESTIMATES

## Examples:

- 1000 tokens  $\approx$  750 words
- Example pricing: \$0.01 per 1K input, \$0.03 per 1K output
- Typical request: 500 input + 500 output =  $\sim$ \$0.02
- Scale: 1000 requests/day =  $\sim$ \$20/day =  $\sim$ \$600/month

## Monitor and adjust:

- 
- 
- 
-

## **Monitor and adjust:**

- Track actual token usage
- 
- 
-

## **Monitor and adjust:**

- Track actual token usage
- Set usage limits
- 
-

## **Monitor and adjust:**

- Track actual token usage
- Set usage limits
- Optimize prompts
-

## **Monitor and adjust:**

- Track actual token usage
- Set usage limits
- Optimize prompts
- Choose models wisely

# SUMMARY

1.

2.

3.

4.

# SUMMARY

1. **Stateful vs. stateless:** Conversations need state management
- 2.
- 3.
- 4.

# SUMMARY

1. **Stateful vs. stateless:** Conversations need state management
2. **Non-determinism:** Control with temperature/top-p
- 3.
- 4.

# SUMMARY

1. **Stateful vs. stateless:** Conversations need state management
2. **Non-determinism:** Control with temperature/top-p
3. **Streaming:** Better UX, handle chunks incrementally
- 4.

# SUMMARY

1. **Stateful vs. stateless:** Conversations need state management
2. **Non-determinism:** Control with temperature/top-p
3. **Streaming:** Better UX, handle chunks incrementally
4. **Cost:** Token-based, monitor usage, optimize prompts

# BRIDGE TO MODULE 8

What we've learned:

- 

What's next:

**Module 8: The Future of AI & Development** — what's coming next.

# BRIDGE TO MODULE 8

**What we've learned:**

- **How** LLM APIs work differently from standard REST

**What's next:**

**Module 8: The Future of AI & Development — what's coming next.**

# QUESTIONS?

*Module 7 — Developing with AI APIs (Condensed)*

