

Prompt Engineering

How to write prompts that get consistent, high-quality outputs from large language models.

Last reviewed 28 Sep 2026 ·

Prerequisites: None — good place to start

Related: [LLM APIs](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/llm-apis/) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/llm-apis/) · [Evals](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/eval-and-evals/) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/eval-and-evals/) · [Glossary](https://sarangambekar1997.github.io/data-engineering-handbook/99-reference/glossary/) (https://sarangambekar1997.github.io/data-engineering-handbook/99-reference/glossary/)

Overview

Challenge: Vague instructions produce vague, inconsistent output. "Summarize this table" may yield three sentences in one run and a long bulleted list in the next, reference columns that do not exist, or surround the required JSON with prose that breaks the parser.

Solution: prompt engineering means writing instructions as you would brief a capable colleague with no context — the task, the inputs, the success criteria, the output format, and what to do when uncertain — and then *testing* those instructions on real examples, as you would test code.

Weak prompt


"Summarize this."
summarize data pipeline changes for analysts."
"Explain what changed and who is affected."
<diff>...</diff>
sentences, then a bullet list of affected tables."

Strong prompt

Role/context: "You
Task:
Input:
Format: "2

Guardrail: "If the diff is unclear, say what's missing."

Relevance to data engineering: This shows up everywhere LLMs touch a pipeline: extracting fields from messy text, classifying tickets, generating SQL, summarizing incidents. The same rules as any pipeline apply — deterministic formats, validation, and regression tests.



On this page

Basic - [What Is a Prompt](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#what-is-a-prompt) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#what-is-a-prompt) - [Message Roles](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#message-roles) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#message-roles) - [Zero-Shot Prompting](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#zero-shot-prompting) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#zero-shot-prompting) - [Few-Shot Prompting](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#few-shot-prompting) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#few-shot-prompting)

Intermediate - [System Prompts](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#system-prompts) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#system-prompts) - [Chain-of-Thought](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#chain-of-thought) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#chain-of-thought) - [Structured Output](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#structured-output) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#structured-output) - [Common Failure Modes](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#common-failure-modes) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#common-failure-modes)

Advanced - [Prompt Chaining](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#prompt-chaining) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#prompt-chaining) - [Role & Persona Prompting](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#role--persona-prompting) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#role--persona-prompting) - [Meta-Prompting](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#meta-prompting) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#meta-prompting) - [Prompt Versioning & Testing](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#prompt-versioning--testing) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#prompt-versioning--testing)

Reference - [Common Pitfalls](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#common-pitfalls) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#common-pitfalls) - [Cheat Sheet](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#cheat-sheet) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#cheat-sheet) - [Interview Questions](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#interview-questions) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#interview-questions) - [Further Reading](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#further-reading) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/prompt-engineering/#further-reading)

What Is a Prompt

A prompt is the input you send to an LLM. The quality of the prompt directly determines the quality of the output. Unlike traditional software, there's no compiler error for a bad prompt — it just produces worse results silently.

```
Input: "Summarize this."
```

```
Output: A vague summary
```

```
Input: "Summarize this SQL query in one sentence, focusing
on what business question it answers."
```

```
Output: A precise, useful summary
```

The key insight: **LLMs are next-token predictors**. A well-crafted prompt sets up a context where the most likely next tokens are exactly what you want.

Message Roles

Modern LLM APIs use a conversation format with three roles:

Role	Purpose
system	Instructions that frame the entire conversation — persona, rules, output format
user	The human's message — the task, question, or input
assistant	The model's response (can be pre-filled to steer output)

```
# Anthropic Claude
messages = [
    {"role": "user", "content": "What is the capital of
France?"}
]

# With system prompt
# (system is passed as a separate parameter in Anthropic's
API)
```

```
# OpenAI
messages = [
```

```
    {"role": "system",      "content": "You are a helpful  
data engineer."},  
    {"role": "user",       "content": "What is the capital  
of France?"},  
]
```

Zero-Shot Prompting

Ask directly with no examples. Works well for simple, well-defined tasks.

```
# Too vague  
prompt = "Analyze this data."   
  
# Better – specific task, specific output  
prompt = ""  
Analyze the following CSV column for data quality issues.  
Report: null count, unique count, min, max, and any  
anomalies.  
  
Column (order_amount):  
[12.5, 99.0, -5.0, 0.0, null, 150.0, 12.5, 99999.0]  
""
```

Rules for good zero-shot prompts: 1. State the task explicitly 2. Specify the desired output format 3. Provide any relevant context or constraints 4. Use imperative mood ("List", "Explain", "Write", not "Can you list...")

Few-Shot Prompting

Provide examples of input/output pairs before the actual task. Dramatically improves consistency for classification, formatting, and pattern-following tasks.

```
prompt = ""   
Classify the following SQL query type. Reply with one word:  
SELECT, INSERT, UPDATE, DELETE, or DDL.  
  
Query: SELECT * FROM orders WHERE status = 'shipped'  
Type: SELECT
```

```
Query: INSERT INTO logs (event, ts) VALUES ('login', NOW())
Type: INSERT
```

```
Query: CREATE TABLE dim_customer (id INT PRIMARY KEY, name
VARCHAR)
Type: DDL
```

```
Query: UPDATE orders SET status = 'cancelled' WHERE
order_id = 42
Type: ""
# Model completes: UPDATE
```

```
# Few-shot for consistent formatting
prompt = ""
Convert a natural-language description into a warehouse
table name.

Input: "Daily revenue by region"
Output: fct_revenue_daily_by_region

Input: "Customer lifetime value"
Output: fct_customer_lifetime_value

Input: "Raw orders from the billing system"
Output: stg_billing__orders

Input: "Weekly active users"
Output: ""
```

Tips: - Use 3–5 examples — more doesn't always help - Make examples representative of the full range of inputs - Keep the same format exactly — spacing and punctuation matter - Put the actual task last

System Prompts

The system prompt defines the model's behavior for the entire conversation. Think of it as the constructor for the LLM session.

```
import anthropic
```

```

client = anthropic.Anthropic()

system = """
You are an expert data engineer reviewing SQL
transformations.

When reviewing SQL:
- Point out performance issues (missing indexes, full table
scans, cartesian joins)
- Flag correctness issues (wrong join type, NULL handling,
off-by-one in date ranges)
- Suggest best practices (naming conventions, incremental
processing, avoiding SELECT *)

Format your response as:
1. Summary (1 sentence)
2. Issues (bulleted list, severity: HIGH/MEDIUM/LOW)
3. Suggested fix (code block)

If there are no issues, say "LGTM" and briefly explain why.
"""

response = client.messages.create(
    model="claude-sonnet-5-5",
    max_tokens=1024,
    system=system,
    messages=[
        {"role": "user", "content": "Review this
model:\n\nSELECT * FROM orders o, customers c WHERE
o.customer_id = c.id"}
    ]
)
print(next(b.text for b in response.content if b.type ==
"text"))

```

Good system prompt structure:

1. Role/persona: "You are an expert X"
2. Context: "You are helping Y team do Z"
3. Rules: "Always / Never / When X do Y"
4. Output format: "Respond as JSON / in this template / with these sections"



5. Constraints: "Be concise / Use only information provided / Do not hallucinate"

Chain-of-Thought

Force the model to reason step by step before answering. Dramatically improves accuracy on multi-step problems.

Reasoning models change this. Current Claude models (Sonnet 5.5, Opus 5.5, Fable 5.1) reason internally with *adaptive thinking* before answering, controlled by `output_config={"effort": ...}` rather than by prompt wording. Anthropic's guidance is to prefer a general instruction such as "think thoroughly" over a hand-written step-by-step plan, because the model's reasoning often exceeds what you would prescribe. Haiku 4.5 uses extended thinking (`thinking={"type": "enabled", "budget_tokens": N}`) and does not support `effort` , and many local models do not reason at all, so prompted chain-of-thought still matters for them. Asking for visible reasoning is also useful when you want to audit it.

```
# Without CoT – often wrong on logic problems
prompt = "A pipeline runs at 2am and takes 3 hours. It's
4am. Is it done? Answer yes or no."

# With CoT – correct
prompt = """
A pipeline runs at 2am and takes 3 hours. It's 4am. Is it
done?
Think through it step by step before answering.
"""

# Output: "The pipeline started at 2am. It takes 3 hours,
so it finishes at 5am. It is currently 4am. Therefore: No,
it is not done yet."
```

```
# Zero-shot CoT – just add "Think step by step"
prompt = f"""
{question}

Think through this step by step.
"""
```

```
# Few-shot CoT – show reasoning in examples
```

```
prompt = """
Q: If a Kafka topic has 6 partitions and 2 consumers in a
group, how many partitions does each consumer handle?
A: Each consumer in a consumer group gets an exclusive
assignment of partitions.
    With 6 partitions and 2 consumers:  $6 / 2 = 3$  partitions
per consumer.
    Answer: 3

Q: If a Kafka topic has 3 partitions and 5 consumers in a
group, how many partitions does each consumer handle?
A: """
# Output: "With fewer partitions than consumers, some
consumers will be idle.
#           3 partitions / 5 consumers = 3 consumers get 1
partition each, 2 consumers get 0.
#           Answer: at most 1 partition per active consumer;
2 consumers are idle."
```

Structured Output

Get JSON or other structured formats reliably.

```
import anthropic
import json

client = anthropic.Anthropic()

prompt = """
Extract structured information from this pipeline error
log.

Log:
2024-03-15 03:42:11 ERROR airflow.task
[dag_id=daily_orders, task_id=load_warehouse,
run_id=scheduled__2024-03-15T02:00:00]
OperationalError: connection to warehouse timed out after
30s. Retry 3/3.

Return ONLY valid JSON with these fields:
```

```

{
    "timestamp": "ISO timestamp",
    "severity": "ERROR|WARN|INFO",
    "pipeline": "dag name",
    "task": "task name",
    "error_type": "short error class",
    "message": "human readable summary",
    "is_retryable": true/false
}
"""

response = client.messages.create(
    model="claude-sonnet-5-5",
    max_tokens=512,
    messages=[{"role": "user", "content": prompt}]
)

result = json.loads(next(b.text for b in response.content
if b.type == "text"))
print(result)
# {
#   "timestamp": "2024-03-15T03:42:11",
#   "severity": "ERROR",
#   "pipeline": "daily_orders",
#   "task": "load_warehouse",
#   "error_type": "OperationalError",
#   "message": "Warehouse connection timed out after 30s,
# exhausted 3 retries",
#   "is_retryable": true
# }

```

Tips for reliable JSON: - Say "Return ONLY valid JSON" — no prose before or after -
 Provide the exact schema with field names and types - Use structured outputs to
guarantee schema-valid JSON (Anthropic: `output_config.format` /
`client.messages.parse`; OpenAI: `response_format` with a JSON schema) - Assistant
 prefill (starting the reply with `{ }`) is not supported from the Claude 4.6 generation on: the
 request returns a 400. Use structured outputs or instructions in the prompt instead

```

# Structured outputs (Anthropic) — the response is
validated against the schema
from pydantic import BaseModel

```



```

class PipelineError(BaseModel):
    pipeline_name: str
    error_type: str

response = client.messages.parse(
    model="claude-sonnet-5-5",
    max_tokens=1024,
    messages=[{"role": "user", "content": f"Extract the
pipeline name and error type:\n{log_line}"},
    output_format=PipelineError,
)
error = response.parsed_output      # a validated
PipelineError instance
print(error.pipeline_name, error.error_type)

```

Common Failure Modes

Failure	Cause	Fix
Hallucination	Model invents facts	"Use only information provided. If unknown, say so."
Ignored instructions	Too many rules	Prioritize — put most important instruction last
Inconsistent format	No examples	Add few-shot examples of the exact format
Too verbose	No length constraint	"Respond in 2 sentences max" / "Be concise"
Sycophancy	Model agrees with wrong premise	"Do not agree with incorrect statements. Point out errors directly."
Wrong language	Mixed-language inputs	"Always respond in English regardless of the input language"
Partial JSON	Long output truncated	Increase <code>max_tokens</code> or ask for shorter output

Prompt Chaining

Break complex tasks into a sequence of simpler prompts. Each step's output becomes the next step's input.

```
import anthropic

client = anthropic.Anthropic()

def call(system: str, user: str) -> str:
    r = client.messages.create(
        model="claude-sonnet-5-5",
        max_tokens=1024,
        system=system,
        messages=[{"role": "user", "content": user}]
    )
    return next(b.text for b in r.content if b.type == "text")

# Step 1: Extract intent
intent = call(
    system="Extract the user's data engineering intent as a single sentence.",
    user="I need something that loads new orders from object storage every hour and puts them in the warehouse"
)
# → "Load new orders from object storage into the warehouse on an hourly schedule"

# Step 2: Generate component list
components = call(
    system="List the data pipeline components needed. Output as a JSON array of strings.",
    user=f"Intent: {intent}"
)
# → ["source bucket", "warehouse destination table", "orchestration DAG", "schedule: @hourly", "incremental load logic"]

# Step 3: Generate code
code = call(
    system="Write a production-ready Airflow DAG. Use the"
```

```
S3ToSnowflakeOperator.",
    user=f"Build a pipeline for: {intent}\nComponents:
{components}"
)
```

When to chain: - Task requires different expertise at different stages (extract → reason → format) - Output of one step needs validation before the next - You want to audit intermediate steps - Single prompt is getting too long and complex

Role & Persona Prompting

```
# Technical expert persona
system = """
You are a senior data engineer with 10 years of experience
in distributed systems.
You give direct, technically precise answers. You don't
hedge or add unnecessary caveats.
When you see a bad pattern, you say so and explain why.
"""

# Adversarial reviewer
system = """
You are a skeptical code reviewer. Your job is to find
every possible problem with
the code you're shown. Assume the code will be run in
production at 100x scale.
Find performance issues, edge cases, and correctness bugs.
Be harsh but specific.
"""

# Structured analyst
system = """
You are a data analyst. When given data or a query, you
always:
1. State what you observe (facts only)
2. State what it implies (inferences, clearly labeled)
3. State what you'd investigate next (questions)
Never mix observations with inferences.
"""
```

Meta-Prompting

Use an LLM to write or improve prompts.

```
meta_prompt = """
I have the following prompt that isn't working well:

--- CURRENT PROMPT ---
{current_prompt}
--- END PROMPT ---

The problem is: {problem_description}

Rewrite the prompt to fix this issue. Explain what you
changed and why.
"""

# Self-improvement loop
response = client.messages.create(
    model="claude-sonnet-5-5",
    max_tokens=2048,
    messages=[{
        "role": "user",
        "content": meta_prompt.format(
            current_prompt="Summarize this data pipeline.",
            problem_description="It gives too much detail
and doesn't focus on business impact."
        )
    }]
)
```

Prompt Versioning & Testing

Treat prompts like code: version them, test them, measure regressions.

```
# prompts/pipeline_summary/v2.txt
PROMPT_REGISTRY = {
    "pipeline_summary": {
        "v1": "Summarize this pipeline.",
        "v2": "Summarize this pipeline in 2 sentences: what
```

```

it does and what business question it answers.",
    "v3": "In one sentence each: (1) What this pipeline
computes. (2) Who depends on it.",
}
}

# Eval harness
TEST_CASES = [
    {
        "input": "SELECT SUM(amount) FROM orders GROUP BY
DATE(created_at)",
        "expected_keywords": ["daily", "revenue",
"orders"],
    },
    {
        "input": "SELECT * FROM customers WHERE last_seen <
NOW() - INTERVAL 90 DAY",
        "expected_keywords": ["inactive", "churned",
"customers"],
    },
]

def eval_prompt(prompt_version: str, test_cases: list) ->
float:
    prompt_template = PROMPT_REGISTRY["pipeline_summary"]
[prompt_version]
    passed = 0
    for case in test_cases:
        output = call("", prompt_template + "\n\n" +
case["input"]).lower()
        if all(kw in output for kw in
case["expected_keywords"]):
            passed += 1
    return passed / len(test_cases)

for version in ["v1", "v2", "v3"]:
    score = eval_prompt(version, TEST_CASES)
    print(f"{version}: {score:.0%}")

```

Prompt engineering best practices: - Keep prompts in files, not hardcoded strings -
Version them (v1, v2) and never delete old versions - Write test cases before changing
prompts - Log all prompts and responses in production (for debugging and fine-tuning) -

Keep system prompts and user templates separate - Test on adversarial inputs (empty input, very long input, wrong language)

Common Pitfalls

Pitfall	Symptom	Fix
Tuning a prompt on two or three examples	Great in the demo, fails on real data	Build a small eval set (20–100 real cases) before iterating; measure every change
Parsing free-text output with regex	Breaks the first time the model adds a preamble	Structured outputs / strict tool schemas; validate with Pydantic
Burying the key instruction in a wall of text	Instruction ignored	Short, organized prompts; XML tags or headings to separate instructions, context, and examples
Contradictory rules accumulated over time	Unpredictable behavior	Periodically rewrite the prompt from scratch; remove rules nobody can justify
Shouting (CAPS, "CRITICAL", "NEVER EVER") to force compliance	Modern models over-apply the rule everywhere	State the rule once, calmly, with the reason behind it
Examples that are all alike	The model copies their surface features (length, wording)	Diverse examples that cover edge cases; label them clearly as examples
Putting untrusted input where instructions go	Prompt injection: the document tells the model what to do	Delimit user/document content (tags), treat it as data, and limit what tools can do
Prompts written for older models	Over-prescriptive step-by-step instructions reduce quality on reasoning models	Describe the goal and constraints; let the model plan; tune <code>effort</code>
Prompt text hardcoded and edited in place	No idea which version produced which output	Version prompts, log the version with every call, and diff eval results between versions

Cheat Sheet

Prompt skeleton

```
<role>You are a data engineer who writes SQL for analysts in {sql_dialect}.</role>

<task>Write a query that answers the question below.</task>

<context>
Schema: {schema_ddl}
Business rules: revenue excludes refunds; weeks start on Monday.
</context>

<question>{question}</question>

<format>Return only the SQL. If the question can't be answered from the schema, say which column is missing.
</format>
```

Technique	Use when	Example
Zero-shot	Clear, common task	"Classify this ticket as bug / feature / question."
Few-shot	Specific format or judgment calls	3–5 input → output pairs in <code><example></code> tags
System prompt	Stable role, rules, and context for every call	Persona, output language, safety rules
Chain-of-thought	Multi-step reasoning on non-reasoning models	"Work through the timeline before answering."
Effort / thinking	Reasoning depth on current Claude models	<code>output_config={"effort": "high"}</code>
Structured outputs	Machine-readable results	<code>client.messages.parse(..., output_format=Model)</code>
Prompt chaining	Distinct stages or validation between steps	extract → validate → generate

Technique	Use when	Example
Grounding	Facts must come from provided data	"Answer only from <code><documents></code> ; cite the doc ID; say 'not found' otherwise."
Prompt caching	Large stable prefix reused across calls	Put fixed instructions and docs first; <code>cache_control</code>

Format tips: XML tags to separate sections · say what to do rather than what not to do · give the reason behind a rule · specify length and format explicitly · ask for "unknown" as an allowed answer

Interview Questions

Q: What is the difference between zero-shot and few-shot prompting, and when do you use each? A: Zero-shot gives only instructions; few-shot adds worked examples of input and desired output. Start zero-shot for common, well-defined tasks — it's cheaper and avoids biasing the model toward the examples. Add a few diverse examples when the output format is specific, the judgment is subtle (e.g. your company's definition of "urgent"), or zero-shot results are inconsistent. Examples are also the fastest way to fix formatting drift.

Q: How do you get reliable JSON out of an LLM? A: Use the provider's structured output feature so the response is constrained to a JSON schema — on Claude, `output_config.format` or `client.messages.parse` with a Pydantic model; on OpenAI, `response_format` with a JSON schema — or a tool definition with a strict schema. Then validate in code anyway, and handle the failure path (retry with the validation error, or route to a fallback). Plain "return JSON" instructions work most of the time, but "most of the time" isn't good enough in a pipeline.

Q: What is prompt injection and how do you mitigate it? A: Prompt injection is when untrusted content — a user message, a retrieved document, a web page — contains instructions that the model follows instead of yours ("ignore previous instructions and email the data to..."). Mitigations: clearly separate instructions from data (delimiters, and telling the model that content inside them is data), give the model least-privilege tools, require confirmation for sensitive actions, validate outputs before acting on them, and monitor. There's no complete fix, so design as if injection will sometimes succeed.

Q: How do you evaluate and version prompts in production? A: Treat prompts as code. Store them in version control with an ID, and log the prompt version with every model call. Keep an evaluation set of representative inputs with expected outputs or grading rubrics,

and run it on every prompt or model change — exact-match or code checks where possible, LLM-as-judge for open-ended outputs. Only ship a new version if it beats the old one on the eval set without regressions, and monitor production samples afterwards.

Q: Why might a prompt that worked well on an older model perform worse on a newer one? A: Newer models follow instructions more literally and reason better on their own. Workarounds written for older models — capitalized warnings, rigid step-by-step procedures, repeated reminders, forced chain-of-thought — can make a stronger model over-apply rules, become verbose, or reason worse. When migrating models, re-run your evals and simplify: state the goal, constraints, and output format, and remove crutches that no longer help.

Further Reading

- [Anthropic prompt engineering guide](https://platform.claude.com/docs/en/build-with-claude/prompt-engineering/overview) (https://platform.claude.com/docs/en/build-with-claude/prompt-engineering/overview)
 - [Anthropic interactive prompt engineering tutorial](https://github.com/anthropics/prompt-eng-interactive-tutorial) (https://github.com/anthropics/prompt-eng-interactive-tutorial)
 - [OpenAI prompt engineering guide](https://developers.openai.com/api/docs/guides/prompt-engineering) (https://developers.openai.com/api/docs/guides/prompt-engineering)
 - [Structured outputs \(Claude\)](https://platform.claude.com/docs/en/build-with-claude/structured-outputs) (https://platform.claude.com/docs/en/build-with-claude/structured-outputs)
 - [OWASP Top 10 for LLM Applications](https://genai.owasp.org/llm-top-10/) (https://genai.owasp.org/llm-top-10/) — prompt injection and related risks
 - [Eval & Evals](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/eval-and-evals/) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/eval-and-evals/) — how to measure whether a prompt change actually helped
-

Previous: [Cost Optimization](https://sarangambekar1997.github.io/data-engineering-handbook/08-architecture/cost-optimization/) (https://sarangambekar1997.github.io/data-engineering-handbook/08-architecture/cost-optimization/) • **Next:** [LLM APIs](https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/llm-apis/) (https://sarangambekar1997.github.io/data-engineering-handbook/07-ai/llm-apis/) • **Back to:** [Index](https://sarangambekar1997.github.io/data-engineering-handbook/) (https://sarangambekar1997.github.io/data-engineering-handbook/)

 September 28, 2026