AI PRODUCT MANAGEMENT 2 · DETAILED COURSE

Prompt Engineering and Context Engineering

Writing and managing the instructions and context that make model behaviour reliable in production.

Chapter 13 of 26 · 23 topics · Edition 2026 · skillnavora.in

CHAPTER 13

Prompt Engineering and Context Engineering

Writing and managing the instructions and context that make model behaviour reliable in production.

Prompt Engineering Fundamentals

Giving clear task, context, constraints and format so the model produces the output you need.

A good prompt states the task, context, constraints and output format. Clear instructions reduce variation and make outputs easier to evaluate.

Prompt Structure

Typical parts: role, task, context, examples, constraints and output format.

Typical parts are role, task, context, examples, constraints and output format. Keep the order consistent so teams can read and edit prompts quickly.

System Prompts

Persistent instructions that set role, rules and behaviour for a whole conversation.

A system prompt sets persistent role, rules and behaviour for a whole conversation or product. It is the first place to put product policy, tone and safety rules.

User Prompts

The per-turn request or input from the user.

The user prompt is the per-turn request. Treat its content as untrusted input, especially when it is combined with company data.

Role-Based Prompting

Assigning a persona or expertise to shape tone and focus.

Asking the model to act as an expert, such as a senior tax advisor, can shape tone and focus. It does not add knowledge the model lacks, so verify the content.

Zero-Shot Prompting

Asking for the task with no examples.

Zero-shot asks for the task with no examples. It is the fastest approach and works well for common tasks with clear instructions.

One-Shot Prompting

Providing a single example of the desired output.

One-shot gives a single example of the desired output. It helps the model match format or style.

Few-Shot Prompting

Providing several examples so the model infers the pattern.

Few-shot provides several examples so the model infers the pattern. Choose examples that cover typical and difficult cases.

Instruction-Based Prompting

Explicit step-by-step instructions and rules for the task.

Explicit step-by-step instructions and rules define the task. Numbered steps and clear do and do-not lists reduce ambiguity.

Chain-of-Thought Concepts

Asking the model to reason through steps, which can improve accuracy on multi-step problems.

Asking the model to reason through steps can improve accuracy on multi-step problems. For products, consider whether the reasoning should be shown to users or kept internal.

Prompt Chaining

Splitting a task into several prompts where each output feeds the next.

A task is split into several prompts where each output feeds the next, such as extract, then validate, then summarise. Chains are easier to test than one very long prompt.

Structured Prompting

Using sections, delimiters and schemas to make prompts unambiguous.

Sections, delimiters and schemas make prompts unambiguous. Use clear markers around user-provided text so the model does not confuse it with instructions.

JSON-Based Outputs

Requesting JSON with defined fields and validating it in code.

Requesting JSON with defined fields lets software use the result. Validate it in code and handle malformed output with a retry or a fallback.

Prompt Templates

Reusable prompts with placeholders filled at run time.

Templates are reusable prompts with placeholders filled at run time. They keep wording consistent across features and make changes easier to review.

Dynamic Prompt Variables

Inserting user, document or session data into templates safely.

Variables insert user, document or session data into a template. Escape or delimit inserted text, so it cannot change the instructions.

Context Engineering

Deciding exactly what information, tools and memory the model sees at each step.

Context engineering decides exactly what information, tools and memory the model sees at each step. It often matters more than the wording of the prompt.

Context Window Optimisation

Prioritising, summarising and trimming content to fit limits while keeping what matters.

Prioritise, summarise and trim content to fit the window while keeping what matters. Put the most important information where the model attends most reliably, usually near the instructions or the question.

Memory and Context Management

Storing and recalling conversation facts, summaries and user preferences across turns and sessions.

Memory stores conversation facts, summaries and preferences across turns and sessions. Ask users what should be remembered and let them delete it.

Prompt Versioning

Tracking prompt changes like code, with versions, owners and test results.

Track prompt changes like code, with versions, owners and test results. Without versioning, teams cannot tell which change caused a quality drop.

Prompt Evaluation

Scoring prompt versions against a fixed test set before release.

Score prompt versions against a fixed test set before release. Use the same inputs for every version so the comparison is fair.

Prompt Injection Awareness

Untrusted text can try to override instructions; treat external content as data, not commands.

Untrusted text, including content from web pages or documents, can try to override instructions. Treat external content as data, not commands, and limit what tools the model can use.

Prompt Guardrails

Input and output checks, refusals and policy rules that keep behaviour within bounds.

Input and output checks, refusals and policy rules keep behaviour within bounds. Guardrails are layers, so combine rules, classifiers and human review for high-risk cases.

Production Prompt Management

Central storage, review, rollout, monitoring and rollback of prompts used in live products.

Production prompts need central storage, review, staged rollout, monitoring and rollback. Treat a prompt change like a code change with the same release discipline.

Summary of terms

Term	Meaning
Prompt Engineering Fundamentals	Giving clear task, context, constraints and format so the model produces the output you need.
Prompt Structure	Typical parts: role, task, context, examples, constraints and output format.
System Prompts	Persistent instructions that set role, rules and behaviour for a whole conversation.
User Prompts	The per-turn request or input from the user.
Role-Based Prompting	Assigning a persona or expertise to shape tone and focus.
Zero-Shot Prompting	Asking for the task with no examples.
One-Shot Prompting	Providing a single example of the desired output.
Few-Shot Prompting	Providing several examples so the model infers the pattern.
Instruction-Based Prompting	Explicit step-by-step instructions and rules for the task.
Chain-of-Thought Concepts	Asking the model to reason through steps, which can improve accuracy on multi-step problems.
Prompt Chaining	Splitting a task into several prompts where each output feeds the next.
Structured Prompting	Using sections, delimiters and schemas to make prompts unambiguous.
JSON-Based Outputs	Requesting JSON with defined fields and validating it in code.
Prompt Templates	Reusable prompts with placeholders filled at run time.
Dynamic Prompt Variables	Inserting user, document or session data into templates safely.
Context Engineering	Deciding exactly what information, tools and memory the model sees at each step.
Context Window Optimisation	Prioritising, summarising and trimming content to fit limits while keeping what matters.
Memory and Context Management	Storing and recalling conversation facts, summaries and user preferences across turns and sessions.
Prompt Versioning	Tracking prompt changes like code, with versions, owners and test results.
Prompt Evaluation	Scoring prompt versions against a fixed test set before release.
Prompt Injection Awareness	Untrusted text can try to override instructions; treat external content as data, not commands.
Prompt Guardrails	Input and output checks, refusals and policy rules that keep behaviour within bounds.
Production Prompt Management	Central storage, review, rollout, monitoring and rollback of prompts used in live products.

Practice exercise: Write a system prompt with a role, rules and JSON output, then version it and test five variations on a fixed set of inputs.

Self-check

1. Explain Prompt Engineering Fundamentals in two sentences to a non-technical stakeholder, then say what you would measure for it.
2. Explain Structured Prompting in two sentences to a non-technical stakeholder, then say what you would measure for it.
3. Explain Production Prompt Management in two sentences to a non-technical stakeholder, then say what you would measure for it.

Related SkillNavora material

- Prompt and Context Engineering : 13 classes with templates.
- AI Skills Academy : Every AI skills course in one place.
- AI Tutor : Ask questions on any topic.
- Study Library : Every lesson, class and PDF by programme.

© SkillNavora · skillnavora.in · Facts and tools change; check official sites before relying on a figure.

ORIGINAL CONTENT

Original AI Product Management content

The AI Product Management programme already on SkillNavora, included unchanged so the full course is in one place.

Quick links Study Library (all material) Learning sites for PM AI builder tools (how-to) Webhooks & Playwright MCP Webhook URLs and keys (start here) AI Skills Academy AI Skills Map Context engineering (PDF) AI skills & templates (PDF) Orchestration (PDF) PRD library (120 ideas) AI PRD Writer AI PM Toolkit AI Evals (PDF) Browser automation (PDF) AI second brain (PDF) PM handbook Top 100 PM questions AI product sense Q&A (86) PM practice test Practice Lab Tool tutorials AI Agents & MCP for PMs

PRD documents

PRD library: template, framing guide and 120 PRD ideas

Learn how top PMs frame a PRD, download the free template, and turn any of 120 ideas across 15 industries into a full PRD with AI.

Open the PRD library ✨ AI PRD Writer PRD template (Word) AI Evals for PMs (PDF)

Context window

Context window: what a model can see at once

The context window is the most tokens a model can consider in one request. It covers your prompt, any documents you paste, the conversation so far and the reply. Anything outside it is invisible to the model. For PMs: a longer window fits more material but costs more and can slow responses; put the most important facts near the instructions; use retrieval (RAG) rather than pasting whole libraries; and test answers on long inputs, where quality can drop.

Context engineering course Context windows in the AI PM 2 course

Tokens and context in the Grokbot course

AI-First Product Management

Resources for Product Managers

The PM Master Handbook, AI and product tools with their official websites, interview practice, a frameworks cheat sheet, a PRD template and free courses.

Open the PM handbook library: all 68 PDFs ↗ PM Glossary Handbook Tokens, Pricing and Tokenization Vibe Coding addendum

PM Master Handbook by SkillNavora
23 modules and 620+ pages: 12 Product Core modules, an 11-class AI Tools Lab, frameworks, case studies, projects and interview prep.

Open the interactive handbook ↗ Complete handbook PDF (620+ pages, 23 MB) ↗

Product Core modules

1Fundamentals of Product Management*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
2Product Strategy and Strategic Thinking*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
3Problem Discovery and Market Research*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
4Problem Framing and Prioritization*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
5UX and Design Collaboration*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
6Effective Communication and Stakeholder Management*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
7Product Analytics and Metrics*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
8Product Launch, Growth, and Monetization*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
9Tech 101 and System Design for PMs*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
10LLMs, RAGs, and AI Foundations*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
11Vibe Coding for Product Managers*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
12Agentic AI and AI Systems*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides

AI Tools Lab

13AI Tools Landscape and Your Starter StackPDF ↗ 14Claude for Daily PM WorkPDF ↗ 15Claude Code: Basic to AdvancedPDF ↗ 16Lovable: Basic to AdvancedPDF ↗ 17Replit and Other AI App BuildersPDF ↗ 18Figma and AI Design ToolsPDF ↗ 19Notion, PRD Tools and AI DocumentationPDF ↗ 20Automation with Zapier, Make and n8nPDF ↗ 21AI Chatbots for BusinessPDF ↗ 22Voice Agents for BusinessesPDF ↗ 23Learning Paths, Courses and CommunitiesPDF ↗

Business AI Stack (chapters 24 to 31)

24The business AI stack: tools and where to learn themPDF ↗ 25Next.js: customer-facing dashboardsPDF ↗ 26Zoho CRM: managing customers and leadsPDF ↗ 27WhatsApp Business Platform: customer messagingPDF ↗ 28Deployment with Vercel and AWSPDF ↗ 29Jira: tracking product and development workPDF ↗ 30Meta Business Suite, social media marketing and Google AdsPDF ↗ 31A one-page website that generates leadsPDF ↗

Complete AI tool guides (15)

1n8n: workflows, Gmail and an AI reply agentPDF ↗ 2Make: scenarios, routers and lead routingPDF ↗ 3Zapier: invoices from Gmail to Drive and a sheetPDF ↗ 4Webhooks: how apps talk in real timePDF ↗ 5Replit: one app for the web and a phonePDF ↗ 6Lovable: a job application trackerPDF ↗ 7Figma: components, auto layout and handoffPDF ↗ 8Claude Code: an AI coding agent in your projectPDF ↗ 9MCP: connecting AI assistants to your toolsPDF ↗ 10LLMs: how they work and how to call themPDF ↗ 11RAG: answers from your own documentsPDF ↗ 12AI agents: a job-search agent with approvalsPDF

↗ 13Agentic AI: planning, guardrails and teamsPDF ↗ 14Generative AI: uses, limits and a weekly reportPDF ↗ 15Other AI tools: choosing one per jobPDF ↗

Complete AI tool guides (one PDF, 30 pages) ↗

Libraries

★Framework Library*29 frameworks with worked examples and templates*PDF ↗ ★Case Study Library*12 product case studies*PDF ↗ ★Practical Projects and Portfolio*5 end-to-end portfolio projects*PDF ↗ ★Interview Handbook*100+ PM interview questions with model answers*PDF ↗ ★AI Tools Lab*Every AI tools class with step-by-step labs*PDF ↗

PM interview practice (free, on the original sites)

Hand-picked public question banks and guides from other PM communities. Each link opens the original page in a new tab.

PM frameworks cheat sheet

Framework	Use it to	How it works
RICE	Prioritise a backlog	Reach × Impact × Confidence ÷ Effort. Score each idea, then sanity-check the ranking.
Kano	Decide which features delight	Classify features as Must-be, Performance, Attractive, Indifferent or Reverse using a short survey.
MoSCoW	Scope a release	Must have, Should have, Could have, Won't have (this time).
Jobs-to-be-done	Understand motivation	When [situation], I want to [motivation], so I can [outcome].
AARRR (Pirate metrics)	Map the growth funnel	Acquisition, Activation, Retention, Referral, Revenue.
HEART	Measure user experience	Happiness, Engagement, Adoption, Retention, Task success, each with goals, signals and metrics.
North Star metric	Align the team	One metric that captures the value users get, supported by input metrics the team can move.
Five whys	Find the root cause	Ask 'why?' repeatedly until you reach a cause you can act on.
CIRCLES	Answer product design questions	Comprehend, Identify users, Report needs, Cut through prioritisation, List solutions, Evaluate trade-offs, Summarise.

PRD template

PRD: [Feature name] Owner: [Name] · Status: Draft · Last updated: [Date] 1. Problem Who has the problem, what it is, and the evidence (data, interviews, tickets). 2. Goal and success metrics Primary metric (with baseline and target), secondary metrics, guardrails. 3. Users and use cases Target segment, top 3 user stories: "As a [user], I want [action] so that [benefit]." 4. Solution What we'll build (link designs), what's out of scope, open questions. 5. AI behaviour (if applicable) Inputs, model, prompt/approach, what it must never do, human review, evaluation set and pass criteria, cost per request. 6. Launch plan Experiment design, rollout stages, support and comms, rollback criteria. 7. Risks and dependencies Technical, legal/privacy, operational; owners and mitigations.

Free courses