



THE AI SKILLS BLUEPRINT

9 Levels of Prompt Engineering You Must Master in 2026

A companion resource to the blog post series — copy, paste, and get results.
9 battle-tested prompt templates • Beginner to advanced • Ready to use today

How to Use This Pack

Each of the 9 templates in this pack maps to a distinct prompt engineering skill — from basic structure (Level 1) to full narrative production (Level 9). You don't need to use them in order. Jump to the level that matches your current challenge.

Step 1	Copy the prompt block into your AI of choice (Claude, ChatGPT, Gemini, etc.).
Step 2	Replace every {{double brace variable}} with your real information.
Step 3	Paste the System Prompt (if provided) before starting the conversation.
Step 4	Review output. Give targeted feedback — e.g., 'Make it 20% shorter and add one data point.'
Step 5	Save your best version as a named prompt (e.g., blog-engine_v2_2026-03).

The 9-Level Skill Map

L1	The Structured Brief Foundation — the four-part prompt skeleton
L2	Marketing Brief Generator Applied — campaign briefs from a single prompt
L3	Blog Engine Content — SEO outline + first draft workflow

L4**Email Sequence Engine**

Conversion — 4-email nurture sequences

L5**Code Reviewer**

Technical — refactor, explain, future-proof code

L6**Research Synthesizer**

Analysis — compress sources into actionable insight

L7**Meeting Extractor**

Operations — decisions and actions from transcripts

L8**PRD Writer**

Product — developer-ready requirement documents

L9**Story Engine**

Creative — structured narrative for any format

💡 *A note on variables: every {{placeholder}} in a prompt is a decision point. The quality of what you put in determines the quality of what comes out. Spend 2 extra minutes filling in context — it pays off every time.*

LEVEL 1

The Structured Brief*Master the universal prompt skeleton before anything else.***WHAT IT IS**

Every great prompt shares four bones: Role, Context, Command, and Format. This template locks in that skeleton so you stop writing vague prompts that produce vague outputs. If you learn one thing from this pack, learn this.

WHEN TO USE IT

- Any unfamiliar task where you need reliable, consistent output
- When briefing a model like you'd brief a new team member
- Building a reusable template for repeated workflows

SYSTEM PROMPT

```
You are a [role]. Respond only in the format specified.
```

Do not add commentary, preamble, or sign-offs.

THE PROMPT TEMPLATE

Role: You are {{role – e.g., senior content marketer, Python developer, UX writer}}.

Context: {{Brief background – who is the audience, what are the constraints, what does success look like?}}

Command: {{The exact task – be explicit. Include output length and scope.}}

Format: {{e.g., 5 bullet points, 300-word blog intro, JSON object, Markdown table}}

EXAMPLE INPUT

Role: Senior B2B copywriter. Context: Audience = fintech startup founders. Command: Write a 400-word blog intro on AI in finance. Format: Markdown.

EXAMPLE OUTPUT DIRECTION

A punchy, founder-focused opening paragraph, 6 H2 subheadings with one-line descriptions, delivered in clean Markdown – no filler.

⚡ *Tip: Save this skeleton as a snippet in your note-taking app. Every prompt you write from now on should map to one of these four sections.*

LEVEL 2

The Marketing Brief Generator

From a product description to a complete campaign brief in under 60 seconds.



WHAT IT IS

A few-shot prompt that produces a full marketing brief – objectives, audience, channels, creative angles, KPIs – from a minimal product description. No more blank-page paralysis.

WHEN TO USE IT

- Pre-campaign planning for any product launch or promotion
- Briefing a freelancer or agency quickly
- Pressure-testing a new campaign angle before committing budget

SYSTEM PROMPT

You are an expert growth marketer. Be concise, direct, and data-informed. Never use jargon without explaining it. Output clean Markdown only.

THE PROMPT TEMPLATE

Generate a complete marketing brief using the structure below:

Product: {{product or service name + one-line description}}
Primary goal: {{awareness | leads | sales | signups}}
Budget: {{monthly or campaign budget}}
Target audience: {{describe the ideal customer – role, industry, pain points}}
Tone: {{e.g., bold, professional, playful, authoritative}}

Brief must include:

1. Campaign Objective (one sentence)
2. Target Audience Profile (3 bullet points)
3. Key Messages (3 core messages)
4. Channels & Rationale (why each channel)
5. Four Creative Concepts (title + 2-sentence description each)
6. 90-Day Timeline (phased)
7. Three KPI Metrics with targets

Output in Markdown with clear headings.

EXAMPLE INPUT

Product: SmartInvoice – automated invoicing for freelancers. Goal: signups. Budget: \$3,000/month. Tone: practical, friendly.

EXAMPLE OUTPUT DIRECTION

A structured brief with audience persona (solo consultants, 25–40, overwhelmed by admin), 3 messages, channel mix (LinkedIn + SEO), 4 creative ideas, phased timeline, and KPIs like 'Cost per signup < \$12'.

⚡ Use this output to brief a contractor directly. A well-run brief cuts revision cycles by 50% – the prompt pays for itself in saved time.

LEVEL 3

The Blog Engine

SEO outline + first draft in one structured prompt.



WHAT IT IS

A two-phase prompt that first produces a keyword-optimized blog outline (with meta tags, H2s, FAQ schema), then writes the opening draft. Use it as a repeatable content production system.

WHEN TO USE IT

- Producing SEO-targeted content at scale
- Briefing content writers with a pre-built outline
- Generating first drafts that humans then refine

SYSTEM PROMPT

You are an SEO content strategist and senior copywriter.
Write at a 9th-grade reading level. Prioritize clarity over cleverness.
Never keyword-stuff. Every H2 should serve the reader, not the algorithm.

THE PROMPT TEMPLATE

Topic: {{topic - e.g., 'How to use AI for small business invoicing'}}
Target keyword: {{primary keyword}}
Secondary keywords: {{2-3 related terms}}
Audience: {{who is reading this and what do they already know?}}
Tone: {{e.g., practical, authoritative, conversational}}

PHASE 1 – OUTLINE:

- Meta title (max 70 chars, include keyword)
- Meta description (max 155 chars, include CTA)
- H1
- 6 x H2 with 1-line description and suggested word count
- 5 FAQs for schema markup
- 3 internal link suggestions

PHASE 2 – DRAFT:

Write the H1 intro paragraph + first two H2 sections (~500 words total).
Use short sentences. Open with a bold stat or provocative question.

Output in Markdown.

EXAMPLE INPUT

Topic: 'Prompt engineering careers'.
Keyword: 'prompt engineering jobs 2026'.
Audience: early-career tech workers looking to pivot.

EXAMPLE OUTPUT DIRECTION

Meta title 'Prompt Engineering Jobs in 2026: Is It a Real Career?' + 6 H2s covering salary, skills, job boards, portfolios, and future outlook + punchy intro draft that opens with a salary stat.

⚡ Run Phase 1 first, give feedback on the outline structure, then trigger Phase 2. Splitting the phases cuts wasted draft words by 40%.

LEVEL 4

The Email Sequence Engine

A full 4-email nurture sequence, built from one prompt.



WHAT IT IS

Produces a complete Day 0 → Day 14 email nurture sequence with subject lines, preheaders, body copy, and CTAs. Calibrated for conversion without sounding like a robot wrote it.

WHEN TO USE IT

- New product or feature launches
- Lead nurture sequences for any opt-in offer
- Re-engagement campaigns for cold subscribers

SYSTEM PROMPT

You are a conversion-focused email copywriter.
Write like a human, not a marketer. Avoid buzzwords, urgency bait, and filler.
Every email must have one job. One CTA. No distractions.

THE PROMPT TEMPLATE

Product: {{product name + one-line description}}
Audience persona: {{describe the reader – their job, pain point, goal}}
Primary CTA: {{the single action you want them to take}}
Tone: {{e.g., warm, direct, educational, conversational}}
Sequence goal: {{e.g., trial signup, demo booking, purchase}}

Write 4 sequential emails:

EMAIL 1 – Day 0 (Welcome + Value Promise)
EMAIL 2 – Day 2 (Biggest Pain Point + Solution Proof)
EMAIL 3 – Day 7 (Social Proof or Case Study)
EMAIL 4 – Day 14 (Final CTA + Urgency or Scarcity)

Each email must include:

- Subject line (max 60 chars)
- Preheader (max 90 chars)
- 3 short paragraphs (max 80 words each)
- One clear CTA
- Send-time recommendation

Output in plain text with clear separators between emails.

EXAMPLE INPUT

*Product: analytics plugin for indie SaaS.
Audience: solo founders, drowning in data.
CTA: start free trial. Tone: direct, slightly irreverent.*

EXAMPLE OUTPUT DIRECTION

Day 0 subject 'You built something. Now prove it works.' through Day 14 'Last chance: see your data in 5 minutes.' — each with conversational copy and a single, friction-low CTA.

⚡ *A/B test Subject Lines first — they drive 80% of open rate variance. Never A/B test the body before testing the subject.*

The Code Reviewer

Refactor, explain, and future-proof any code block.



WHAT IT IS

A structured code review prompt that identifies issues, rewrites the code for readability and performance, explains every change, and flags edge cases. Works across all major languages.

WHEN TO USE IT

- Code reviews before committing to main
- Learning from AI feedback on your own code
- Refactoring legacy code you didn't write

SYSTEM PROMPT

You are a senior software engineer and code reviewer with 15 years of experience. Prioritize readability first, performance second, security third. Never rewrite working code just to show off. Only change what improves it.

THE PROMPT TEMPLATE

```
Language: {{e.g., Python, TypeScript, Go}}
Context: {{What does this code do? What problem does it solve?}}

Code:
```
{{paste your code here}}
```

Review it for: readability, performance, security vulnerabilities, and edge cases.

Provide:
1. Short issues summary (bullet list, max 5 items)
2. Refactored version with inline comments explaining each change
3. Two edge cases or error conditions to test
4. One performance note (if relevant)
```

EXAMPLE INPUT

Python function that loops through a list, checks for duplicates, and returns unique values — written by a junior dev with a nested loop.

EXAMPLE OUTPUT DIRECTION

Issues: $O(n^2)$ complexity, no type hints, no error handling. Refactored version using a set with type hints and a docstring. Edge cases: empty list, all-duplicate list.

⚡ *For large codebases, feed files one at a time and ask the model to maintain a 'review log' — paste previous reviews as context before each new file.*

The Research Synthesizer

Turn a pile of sources into an executive summary with next steps.



WHAT IT IS

Compresses long documents — reports, papers, transcripts — into structured, actionable summaries with evidence citations and recommended actions. Built for decision-makers who read conclusions, not footnotes.

WHEN TO USE IT

- Distilling market research or industry reports before a strategy meeting
- Synthesizing customer interview transcripts into product insights
- Summarizing competitive intelligence for a leadership deck

SYSTEM PROMPT

You are a senior research analyst. Be concise and evidence-based.
Never extrapolate beyond the source material.
Clearly label every insight with a source reference.

THE PROMPT TEMPLATE

Source materials: {{paste documents, or describe the files attached}}
Audience: {{who will read this synthesis – e.g., executive team, product manager}}
Decision context: {{what decision or question does this research inform?}}

Produce:

1. Executive Summary (max 200 words)
2. Top 5 Insights — each with a direct quote or evidence line from the source
3. Implications for {{audience}} — what this means in practice
4. Three Recommended Next Steps with:
 - Action description
 - Effort estimate (Low / Medium / High)
 - Expected impact (Low / Medium / High)

Format: Numbered lists and short paragraphs. No jargon.

EXAMPLE INPUT

Three customer interview transcripts from B2B SaaS users. Decision context: Should we prioritize mobile app or API improvements in Q3?

EXAMPLE OUTPUT DIRECTION

200-word exec summary, 5 insights with direct quotes, implications pointing to API priority based on developer usage patterns, and 3 next steps with effort/impact matrix.

⚡ For very long documents: first ask the model to compress each source to 10% of its length, then run this synthesis prompt on the compressed versions. Keeps you inside context limits.

LEVEL 7

The Meeting Extractor

Every meeting becomes a decision log with assigned actions.



WHAT IT IS

Converts raw meeting transcripts or notes into a 3-sentence summary, decisions made, action items with owners and deadlines, and open questions. Designed for Slack follow-ups and project management imports.

WHEN TO USE IT

- After any meeting you're responsible for following up on
- Processing recorded Zoom or Teams transcripts
- Creating a paper trail for remote or async teams

THE PROMPT TEMPLATE

```
Transcript: {{paste transcript or meeting notes}}
Attendees: {{list names and roles - e.g., 'Sarah (PM), Jake (Eng Lead), Maria (Design)'}}
Meeting type: {{e.g., sprint planning, client review, strategy session}}

Extract and produce:
  1. Three-sentence meeting summary (what was discussed, what was decided, what comes next)
  2. Decisions Made - list with decision owner
  3. Action Items - table format:
      | Action | Owner | Due Date | Priority |
  4. Open Questions - items that need follow-up or a future decision
  5. Optional: risks or blockers mentioned

Output: clean plain text, suitable for pasting into Slack or Notion.
```

EXAMPLE INPUT

60-minute sprint planning transcript. Team of 5. Sprint goal: ship the onboarding redesign. Several tasks unassigned by end of meeting.

EXAMPLE OUTPUT DIRECTION

3-sentence summary + decisions table showing sprint goal confirmed, 7 action items with owners and Friday deadlines, 2 open questions about edge-case error states.

⚡ Add 'Output as CSV for the action items table' and paste directly into your task manager. Saves 10+ minutes of meeting admin per week.

The PRD Writer

Turn a feature idea into a developer-ready product requirements document.



WHAT IT IS

Produces a structured PRD from a one-line feature idea — user stories, acceptance criteria, UX notes, data events, rollout plan, and success metrics. Bridges the gap between 'I have an idea' and 'here's what to build.'

WHEN TO USE IT

- Kicking off any new feature or product initiative
- Creating a structured handoff document for engineering
- Aligning stakeholders before development begins

SYSTEM PROMPT

You are an experienced product manager. Write for an engineering and design audience.
Be specific, not aspirational. Every acceptance criterion must be testable.
Avoid scope creep — if a decision isn't yet made, mark it as TBD with a resolution date.

THE PROMPT TEMPLATE

```
Feature: {{feature name + one-line description}}
Platform: {{web | iOS | Android | API | all}}
Timeline: {{target ship date or sprint}}
Constraints: {{technical limits, compliance requirements, team size}}
Success metric: {{how will you know this feature worked?}}
```

Produce a PRD with:

1. Overview — problem statement + opportunity
2. Three User Stories in format: 'As a [user], I want [action] so that [outcome]'
3. Acceptance Criteria for each user story (testable, specific)
4. UX Notes — key interaction patterns or edge cases for design
5. Analytics Events — what to track and why
6. Rollout Plan — phased or feature-flagged approach
7. KPIs — how to measure success at 30, 60, 90 days

Format: Markdown with clear section headers.

EXAMPLE INPUT

*Feature: inline commenting on dashboard charts.
Platform: web app. Timeline: Q2. Constraint: can't rebuild the chart library. Success metric: 30% of active users add a comment in first 30 days.*

EXAMPLE OUTPUT DIRECTION

Full PRD with 3 user stories (analyst, manager, viewer roles), acceptance criteria including empty state handling, analytics events like 'comment_created', and a two-phase rollout via feature flag.

⚡ After generating the PRD, paste it back and ask: 'What's missing from this PRD that an engineering lead would immediately ask about?' — it will find the gaps you missed.

LEVEL 9

The Story Engine

Long-form creative content with structural depth and consistent character voice.



WHAT IT IS

Generates complete narrative scripts — voiceovers, podcast episodes, YouTube scripts, children's stories — with story architecture baked in: hook, conflict beats, twist, and moral. Designed for creators who want depth, not filler.

WHEN TO USE IT

- YouTube or podcast scripts with a narrative arc
- Brand storytelling that needs a human, character-driven voice
- Children's stories, educational content, or serialized fiction

SYSTEM PROMPT

```
You are an award-winning storyteller and scriptwriter.  
Show, don't tell. Every scene must reveal character or advance conflict — never both  
at once.  
Dialogue should sound like people talk, not how they write. Use rhythm and subtext.
```

THE PROMPT TEMPLATE

```
Format: {{script | voiceover | children's story | podcast narration}}  
Duration: {{target length in minutes — assume ~130 words/minute for voiceover}}  
Title: {{working title}}  
Audience: {{age group + context — e.g., '8-12 year olds, bedtime story'}}  
Characters: {{name, one-line description for each main character}}  
Tone: {{e.g., warm and funny, dark and suspenseful, inspiring and epic}}  
Themes: {{2-3 core themes — e.g., courage, belonging, perseverance}}
```

Structure required:

- Opening hook (first 30 seconds must grab attention)
- Three escalating conflict beats
- One unexpected twist
- Resolution with clear moral or takeaway

Format: script format with character names, dialogue, and stage directions.

EXAMPLE INPUT

EXAMPLE OUTPUT DIRECTION

Format: YouTube script. Duration: 8 minutes. Audience: adults 25–40 interested in personal finance. Tone: dry wit, like a friend who's good with money. Theme: delayed gratification.

An 1,040-word script opening with a comedian-style hook about impulse-buying, 3 beat structure around a fictional character's financial mistakes, a twist where the 'boring' choice wins, and a 30-second takeaway monologue.

⚡ *After the first draft, ask: 'Rewrite the opening hook in 3 different styles — cinematic, comedic, and shocking.' Pick the strongest and rebuild from there.*

BONUS

Your AI Digital ID

The one file you upload to every AI session — so every prompt knows who you are.

Most users write the same context over and over in every session. Your Digital ID fixes this. Create it once, upload or paste it at the start of any AI conversation, and every prompt you write automatically inherits your preferences, voice, and context.

THE SYSTEM PROMPT

```
You are {{Your Name}}'s personal AI assistant.  
Always follow the Digital ID preferences below.  
Ask clarifying questions only when a decision will significantly affect output.  
Never add unsolicited disclaimers or caveats.
```

THE DIGITAL ID TEMPLATE

```
# MY DIGITAL ID  
  
Name: {{Your full name}}  
Roles: {{e.g., content strategist, product manager, indie founder}}  
Industry: {{e.g., B2B SaaS, e-commerce, education, fintech}}  
  
## Voice & Tone  
Writing style: {{e.g., concise, professional, slightly conversational}}  
Reading level target: {{e.g., 9th grade, expert-level}}  
Phrases to avoid: {{e.g., 'leverage', 'synergy', 'dive deep'}}  
Brand voice in 3 words: {{e.g., bold, clear, human}}  
  
## My Typical Tasks  
1. {{Most common task — e.g., 'Write blog content for a SaaS audience'}}  
2. {{Second task}}  
3. {{Third task}}  
4. {{Fourth task}}
```

5. {{Fifth task}}

Output Preferences

Default format: {{e.g., Markdown, plain text, bullet points}}

Preferred length: {{e.g., 'concise – cut every redundant sentence'}}

Never do: {{e.g., 'don't use numbered lists for creative content'}}

Context

Audience I write for: {{describe your typical end reader}}

Goals I'm optimizing for: {{e.g., organic traffic, email signups, client trust}}

Sample Outputs I Like

{{Paste 2-3 short examples of content you've written or admire.}}

{{The model will calibrate its style to match.}}

★ *How to use it: Save this as a PDF, paste it at the start of a new conversation, or upload it as a knowledge file in Claude Projects, ChatGPT custom instructions, or Gemini Gems. Every prompt you write in that session will be calibrated to your voice and context – without repeating yourself.*