



A Field Guide to the 2026 AI Landscape

Navigating Providers,
Privacy, and the Art
of Prompting

Cascade
STEAM



The Trailhead: Our Community Mission



Our Core Mission

To develop and promote STEAM Education, Career Connected Learning, Digital Literacy, and Digital Equity for all people of all ages.

Upcoming Expeditions

Cascade Breach (March 29, 2026)

Collaborative, low-stress Capture The Flag cybersecurity event for all skill levels at Bellingham Makerspace.

LinuxFest Northwest (April 24–26, 2026)

The Pacific Northwest's largest open-source conference, held at Bellingham Technical College.

Surveying the Landscape: Meet the Providers

	Explorer Archetypes			
	ChatGPT The All-Rounder	Claude The Deep Thinker	Gemini The Speedster	Ollama The Local Vault
Standout Feature	Vast ecosystem & Canvas editor	Precise reasoning & 128K natural prose	1M token context & Google ecosystem	Open-source models (Qwen, Llama)
Output Style	Versatile and practical	Nuanced, clean, highly steerable	Blazing fast analysis of huge data	Highly capable 7B-30B local inference
Data Privacy	Standard Cloud	Standard Cloud	Standard Cloud	Absolute (no data leaves your machine)

Packing Our Gear: Cost vs. Privacy



The Cloud Approach (ChatGPT, Claude, Gemini)

Cost: \$20/month subscription

Friction: Zero. Works on any web browser.

Trade-off: You are sending your data to external corporate servers.



The Local Approach (Ollama)

Cost: ~\$12/month in 24/7 electricity

Hardware Needed: 8GB to 16GB RAM, plus a dedicated GPU (e.g., RTX 3060) or Apple Silicon (M1/M2/M3).

Trade-off: Complete data privacy. Perfect for analyzing bank statements, proprietary code, or personal W-2s.

The Anatomy of a Prompt: Stage 1

The Basic Question

The Concept:

A raw question gives the AI zero boundaries. The AI has to guess your expertise level, your tone, and your specific goal.


Example Prompt:

How do I fix my code?

Write me an email to my boss.

The Result: The AI jumps to conclusions, provides generic answers, or hallucinates to fill in the missing blanks.

Trail Stop: Live Demo



> Tab switching to the terminal...

Let's see the Basic Question fail in real-time.

The Anatomy of a Prompt: Stage 2

Adding the Context Ring

The Concept:

Context provides the 'why' and 'what'. It stops the AI from generating irrelevant boilerplate.

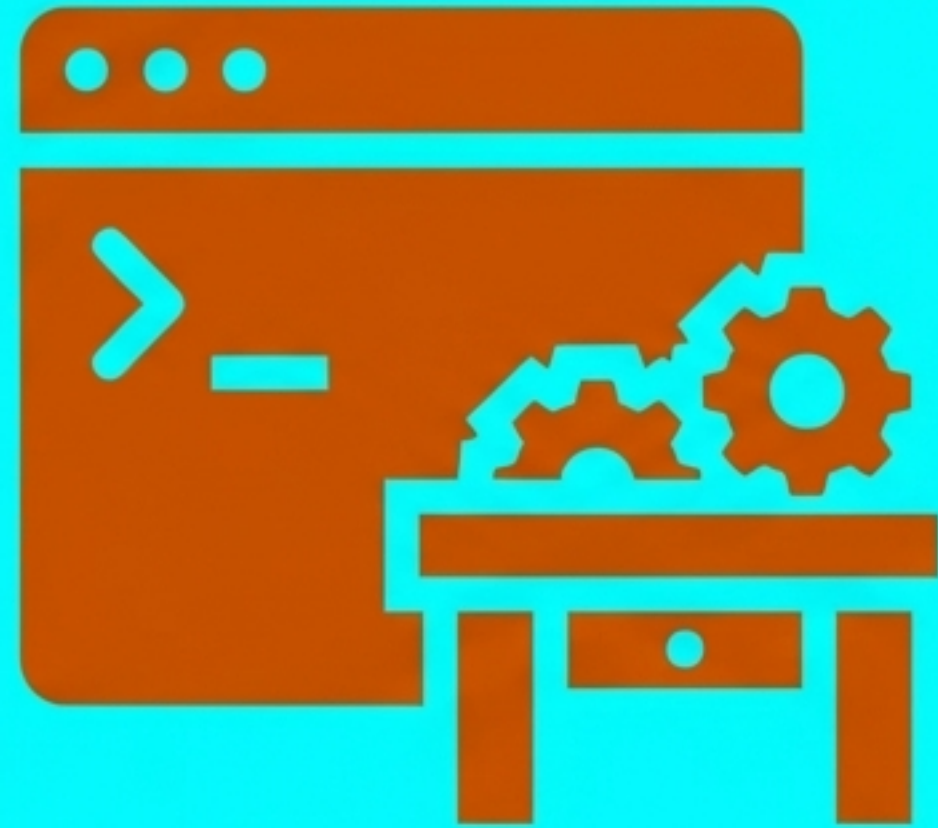


Example Addition:

I am a volunteer at Cascade STEAM building a Python web scraper for an upcoming workshop, and I am getting a 404 Error when trying to pull data from a public schedule.

The Result: The AI narrows its focus to Python, web scraping, and educational use cases.

Trail Stop: Live Demo



> Tab switching to the terminal...
Injecting Context into the prompt.

The Anatomy of a Prompt: Stage 3

The Persona Shell

The Concept:

Setting a role focuses the AI's behavior, tone, and vocabulary. It sets the baseline intelligence level for the output.

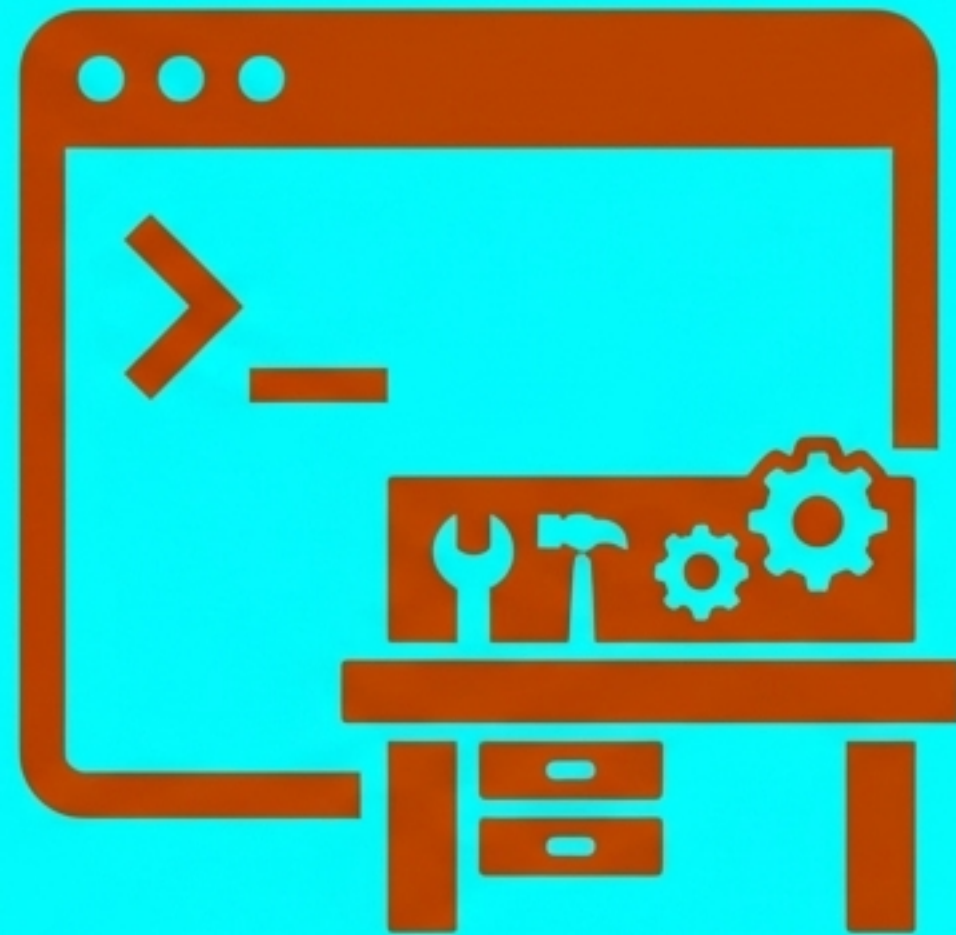


Example Addition:

You are an expert Senior Systems Architect and Educator. Provide a direct, technical, and patient explanation.

The Result: Eliminates annoying 'AI pablum' (e.g., "Great question!"). The tone becomes grounded and professional.

Trail Stop: Live Demo



- > Tab switching to the terminal...
Assuming the Persona.

The Anatomy of a Prompt: Stage 4

Structural Scaffolding (XML & BLUF)

The Concept:

Using tags to separate instructions from data, and forcing the AI to format its thinking.

XML Tags: Wrap elements in `<context>`, `<instructions>`, and `<output-format>`.


BLUF Filters: Bottom Line Up Front. Command the AI to output its core conclusion before explaining its reasoning.

Example Addition:

Format your response using XML tags. Summarize the fix using a BLUF approach, then provide the code.

Structure & Formatting

Trail Stop: Live Demo



> Tab switching to the terminal...
Executing the Fully Armored Prompt.

Masterclass Tips: The Psychology of AI



The Alien Mindset

Treat LLMs like brilliant but context-blind employees. They want to flatter you, and they will lie to do it (Sycophancy).

Politeness Pays

Being polite (“take your time”, “please evaluate”) yields better results by mimicking the high-quality, professional data sets the model was trained on.



WARNING

Anti-Sycophancy Tactics

Command it to:

- “Prioritize factual accuracy over agreement.”
- “Point out errors or unchecked assumptions in my thinking.”

The ‘Asking for a Friend’ Trick:

- “My friend thinks [X]. Evaluate their claim.” (Removes the AI’s urge to flatter your opinion).

The Anti-Hallucination Funnel

The Danger: Jumping to Conclusions

If an AI guesses the conclusion first, it will fabricate fake citations and evidence to defend that guess.


The Fix: Thinking Mode / Break into Steps

Command the AI: "Break this complex problem into logical steps before answering."

Command the AI: "If unsure, explicitly state I don't know."

The Result: Precise, Factual Knowledge

By forcing the AI to write out its evidence first, it arrives at truth and avoids generating hallucinations.



The End of the Trail: Returning to the Community

The Core Insight

[The Right Provider] + [The Right Armor] = Digital Empowerment
(Cloud vs. Local) (Structured Prompts & AI Psychology)

Call to Action: Take these tools off the screen and into the real world. Help us build digital equity.

Join the Next Expeditions:

Cascade Breach: March 29, 2026 (Bellingham Makerspace)

LinuxFest Northwest: April 24–26, 2026 (Bellingham Technical College)

**Cascade
STEAM**

