Prompt Engineering

How to Roll Out Microsoft 365 Copilot Without Restricted SharePoint Search: A CTO's Guide for Swiss SMEs

November 6, 2026

What do you do when the containment feature your Copilot pilot was built on stops accepting new tenants?...

Wells Fargo Ships Section 15 on November 24: Every AI Agent Instruction Is Deemed Authorized by You

October 31, 2026

On October 3, 2026, Wells Fargo posted a revised Online Access Agreement. Buried in it is Section 15:...

SGLang 0.5.19 in Practice: Beam Search, a Faster Cold Start, and What Still Breaks

September 10, 2026

SGLang's 0.5.19 headline is beam search, and it will not run on a disaggregated deployment with speculative decoding....

Anthropic Cuts Fable 5.1 Cache Reads 75% to $0.25/M — and Breaks Three Prompt Patterns Developers Rely On

September 3, 2026

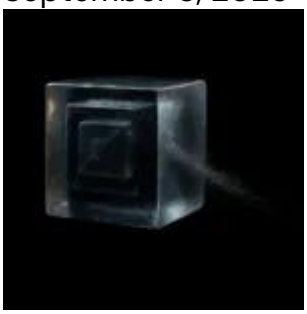

Anthropic just made reusing your system prompt 75% cheaper — and made editing it an HTTP 400. On...

The Harness, Explained: Why the Same Scaffold Scores 65% or 74% on SWE-bench

September 1, 2026

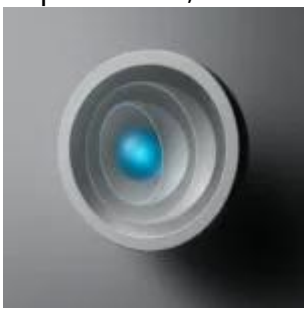

A 100-line Python harness scored 65% on SWE-bench Verified in July 2025. The same project now claims above…

StackOne Defender 0.8.2 in Practice: An Apache-2.0 Prompt-Injection Filter for Tool Calls — What Works, What Does Not

August 30, 2026

A bundled ONNX classifier now sits between your agent and its tools, catching a vendor-reported 88.7% of prompt…

Grok Still Leaks Full Chat Histories 11 Weeks After Disclosure—AES-256 Encrypted Prompts Beat Guardrails 40% of the Time

August 27, 2026

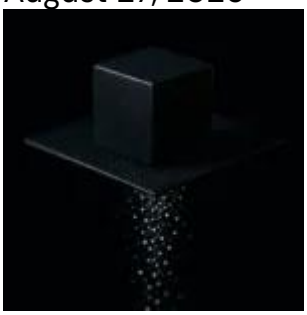

Grok's safety classifier reads your prompt. It does not read what Grok decrypts afterwards. Adversa AI shipped attack…

Apiiro CLI Launches April 9 Turning AI Coding Assistants Into Security Engineers by Rewriting Developer Prompts in Real Time

April 19, 2026

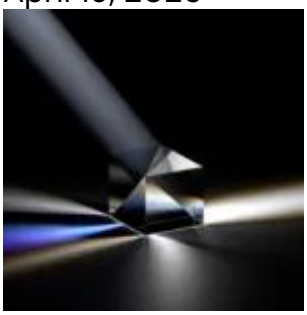

Your AI coding assistant now has a shadow editor—one that rewrites your prompts before the AI model sees…

PromptEng 2025 Workshop Accepts 7 Research Papers on January 24—First Academic Forum Dedicated to Prompt Engineering Techniques

March 23, 2026

While engineers debated whether prompt engineering was “real engineering,” seven research teams quietly got their papers peer-reviewed and…

Handshake58 Launches on Bittensor Subnet 58: $0.0001 Micropayments for AI Agents Using Polygon USDC and Zero-Gas Vouchers

March 12, 2026

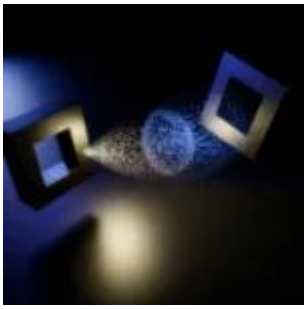

The AI industry spent 2025 building agents that can reason, plan, and act—but forgot to give them wallets.…

xAI Launches Grok 4.20 Beta with Four-Agent Architecture—65% Hallucination Reduction Shifts Prompt Engineering From Iterative Chat to Structured Contracts

February 23, 2026

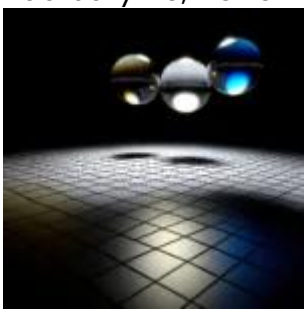

The iterative prompt refinement loop that defined AI workflows for three years just became a liability. Grok 4.20’s…

Calibrated Confidence Prompting: The Silent Shift from Asking Better to Trusting Smarter

January 7, 2026

Everyone’s arguing prompt engineering is dead. Meanwhile, engineers who actually ship LLMs to production discovered something more valuable:…

The Prompt Engineering Paradox: Why Reasoning Models Like o1 and DeepSeek R1 Are Making Traditional Prompt Optimization Obsolete—And What Replaces It

December 16, 2025

Your entire prompt engineering stack just became technical debt. The skills that made you invaluable six months ago...

Mastering Autonomous and Multimodal Prompt Engineering: The New Frontier for Precision AI Workflows in 2025

October 14, 2025

Stop using yesterday's prompt tactics—2025's AI teams are outpacing you with adaptive, autonomous, and multimodal engineering, redefining what's...

Harnessing Automated Multimodal Prompt Engineering for Reliable, Domain-Specific AI Workflows in 2025

October 9, 2025

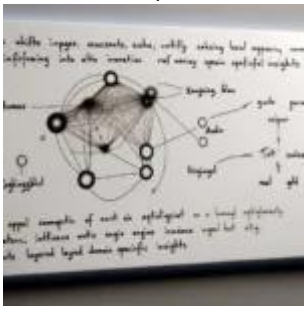

Do you really know what your AI is thinking, or is it just following blind instructions? Next-gen prompt...

Why Mastering Multimodal AI and Prompt Engineering is the New Frontier in AI Art Collaboration for 2025

October 9, 2025

Think AI art tools are just for generating pretty pictures? That mindset is officially obsolete—the real paradigm shift...

Mastering Automated and Multimodal Prompt Engineering: The New Core Competency for 2025 AI Product Teams

September 11, 2025

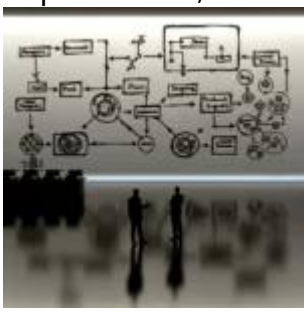

It's not just about crafting clever prompts anymore—the real winners are automating and scaling this discipline before their...

The New Frontier of Concept-Level Control in Multimodal AI: Precise Unlearning Without Full Retraining

September 10, 2025

What if you had the power to erase harmful AI biases and outdated knowledge with surgical precision—no retraining,...

Mastering GPT-5's 'Thinking Mode': Elevating AI-Driven Developer Workflows Beyond Simple Code Generation

September 8, 2025

What if your AI coding co-pilot could finally hold its own in multi-hour whiteboard debates, refactor your legacy...

Harnessing AI-Driven Adaptive Prompt Engineering to Optimize Real-Time Developer Workflows

September 5, 2025

What if your static AI prompts are the silent productivity killer nobody talks about? The future of developer...

[The Hidden Pull Request Tax: Why AI Coding Assistants Are Creating a 91% Review Bottleneck That's Killing Enterprise Productivity]

August 11, 2025

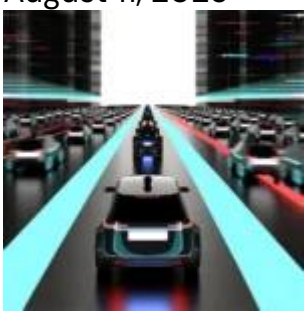

Your AI-powered developers are shipping code at warp speed while your company grinds to a halt. The numbers...