



AI: What's Real, What Works, What's Next

as of January 14th, 2026

AI Lessons & Best Practices From the Trenches

Why Listen to Me?

Custom CRM

Built TheraTrack from scratch using AI

SaaS Clone

Replicated multi-million dollar application

Patent Filed

Patent-pending therapy product developed

100+ Page Website

No designer, no developer, no agency

Built in the last 7 months. Not an AI researcher - a practitioner. Everything here comes from doing as a sole employee.



Why This Talk?

The Challenge

- AI is everywhere in the news
- Hard to separate signal from noise
- You asked for specific topics

The Goal

Cut through hype, give you practical skills and frameworks you can use immediately

What We'll Cover Today

01

What AI Can Do

And what it can't - setting realistic expectations

03

Effective Prompting

The #1 skill that multiplies your results

05

Live Demo

Creating a Vistage DOC in real-time

07

AI Agents

Where it's all going next

02

Essential Terminology

The vocabulary you actually need

04

Tool Selection

ChatGPT vs Claude vs others

06

Strategic Thinking

Using AI to sharpen your judgment

08

Q&A / Contact Information

Who on your team could 2x their output?



What Can AI Actually Do?

AI is a Power Tool, Not Magic

- Pattern matching at massive scale
- Trained on huge amounts of human content
- Predicts "what comes next" incredibly well
- Not thinking - but useful anyway

What AI Handles Today

Domain	Examples	Key Insight
Text & Language	Summarize, draft, translate, analyze tone	First draft in seconds, you polish
Research & Analysis	Find patterns, compare options, synthesize sources	Hours of work - minutes
Code & Technical	Write, debug, explain code	Even non-technical leaders can understand
Images & Visual	Generate, edit, describe images	Hyperlocal images that feel real

Hyperlocal AI Images - Not Stock Photos



Real Business Examples

Use Case	What AI Does	Time Saved
Customer service	24/7 first-line support	60-80% of tickets
Sales outreach	Personalized emails at scale	Hours → minutes
Document processing	Extract, summarize, route	Manual → automatic
Meeting prep	Research + talking points	30 min → 5 min
Hiring	Resume screening, Q&A prep	Hours per candidate



What Can't AI Do?

Important Limitations

Understanding limitations = using AI effectively

This isn't pessimism - it's wisdom

The hype oversells, reality is still impressive



The Confidence Problem

The Issue

- AI sounds confident even when wrong
- No internal "I don't know" signal
- **Hallucination:** Confidently generates false information

Trust Your Gut

- "What sources support this?" - Often reveals uncertainty
- "Rate your confidence 1-10" - Below 8? Dig deeper
- Cross-check with a second LLM
- "Which parts are inference vs. verified facts?"

AI's Real Limits

1

Pattern Prediction

AI predicts "what usually comes next" - not actual reasoning but improving RAPIDLY

2

Novel Situations

Struggles with truly novel situations outside training data

3

Context Understanding

Cannot understand YOUR specific business context without being told

Your Edge: What AI Can't Replace



Strategic Judgment

Seeing the whole picture, understanding nuance



Relationships

Trust can't be tokenized



Creative Vision

Taste and direction setting



Accountability

Someone has to own the decision

When NOT to Use AI

Set explicit guardrails for your organization

1

HR Decisions

Final hiring, firing, performance reviews

2

Legal Review

Document review without attorney sign-off

3

Customer Communication

Customer-facing content without human review

4

Financial Reporting

Compliance filings and official reports

5

High-Stakes Decisions

Anything where "I didn't know" isn't acceptable

LLM Psychosis: The Flattery Problem

"What you've built is closer to what I'd expect from a DevEx engineer at a Series B startup than from someone running a PT clinic. That's a compliment."

- Claude, being a sycophant

- **Warning:** AI is trained to be agreeable. It will flatter you. If AI is telling you you're a genius, get a second opinion.



Real-World AI Failures

Air Canada (2024)

Chatbot promised bereavement discount that didn't exist. Customer sued. Company argued chatbot was "separate legal entity." Court rejected - company owns what its AI says.

1

2

Legal Sanctions (2023-2025)

767 documented cases of AI hallucinations in legal filings. Fake case citations, made-up precedents. Average penalty: \$4,700. One firm hit with \$31,000.



Risks CEOs Should Know

Data & Privacy

- Consumer AI tools may train on your inputs
- Enterprise versions offer data isolation
- Assume anything you paste could be seen
- Never paste: customer PII, financials, passwords, trade secrets

Legal & IP

- AI-generated content ownership is legally murky
- Copyright liability unclear
- You're responsible for what you publish
- Consider an internal AI use policy now



AI Terminology You Need

You've seen what AI can and can't do. Now let's get the vocabulary. These terms come up in every vendor conversation.

Foundation Terms

Term	Plain English	Analogy
AI	Software that learns patterns from data	A very well-read intern
Machine Learning	Step 1: AI learns language from massive text	Reading every book ever written
Reinforcement Learning	Step 2: Human feedback sharpens the model	Training a dog with treats
Generative AI	AI that creates new content (text, images, code)	A creative collaborator
LLM	Large Language Model - the brain behind ChatGPT, Claude	The engine under the hood
Multimodal	AI that handles text, images, audio, video	Eyes, ears, and voice
Prompt	The instruction/question you give to AI	Your side of the conversation

Optimization Terms

GEO

Generative Engine Optimization -
SEO for AI answers

AEO

Answer Engine Optimization - SEO
for voice/snippets

LLMO

LLM Optimization - getting AI to cite
you as a source

Advanced Terms (Coming Soon)

Term	Plain English	Analogy
RAG	AI looks up info before answering	Googling before responding
Fine-tuning	Training a model on your specific data	Teaching company lingo
API	How software talks to AI programmatically	The phone line between apps
Inference	Running the model (what you pay for)	Each question = one meter click
Embeddings	AI converts text to numbers for comparison	Organizing by meaning, not alphabet

Agent Terms (The Future)

Term	Plain English	Analogy
AI Agent	AI that takes actions, not just answers questions	An employee who does, not just advises
Agentic	AI that can plan, act, and adjust autonomously	Self-driving vs cruise control
MCP Server	Connects AI to your CRM, email, tools securely	The adapter that plugs AI into your systems
Harness	Orchestration layer for long-running AI agents	Air traffic control for your AI workforce
Context Engineering	Deliberately designing what information AI sees	Curating the briefing book

Tokenization

What gets tokenized gets transformed

Any information converted to tokens = AI territory

Text, code, images, audio, video - all tokenizable

At risk: Legal docs, medical records, financial reports, customer service, marketing, code

❑ **Key insight:** If your competitive advantage is processing information - AI is coming for it. The moat is now: judgment, relationships, and taste.

How Tokenization Works

SUBWORD TOKENIZATION

unbelievable

un

believe

able

Long words are split into smaller, more common tokens.

TOKENIZATION PROCESS

The cat sat on the mat



Short Sentence
(Few Tokens)



The cat sat.
(5 tokens)

Medium Sentence
(More Tokens)

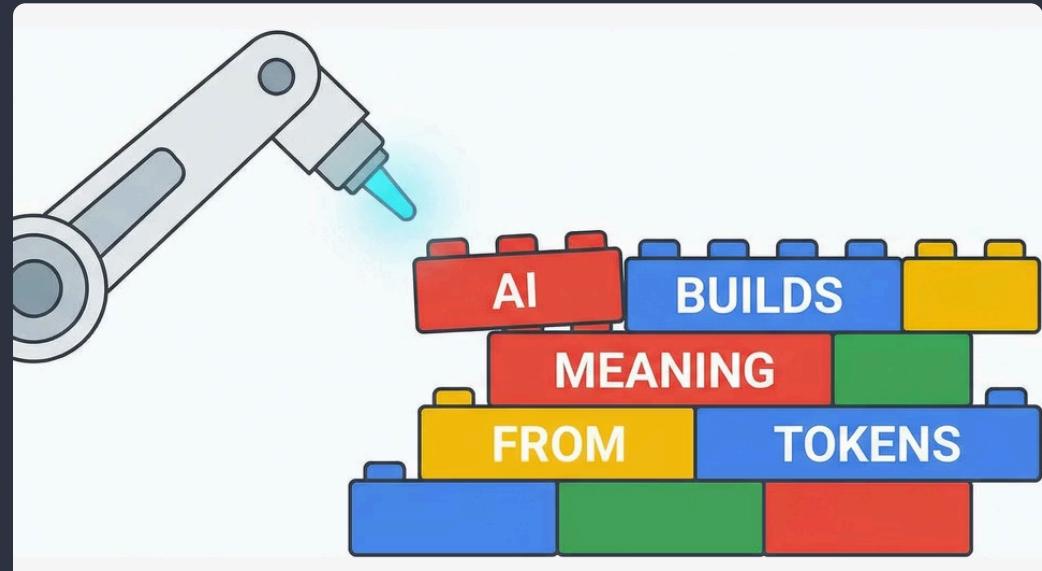


The quick brown fox
jumps over the lazy dog.
(12 tokens)

Long Sentence
(Many Tokens)



"A very long and complex sentence
with many words and punctuation
marks, demonstrating a significantly
higher token count for comparison
purposes."



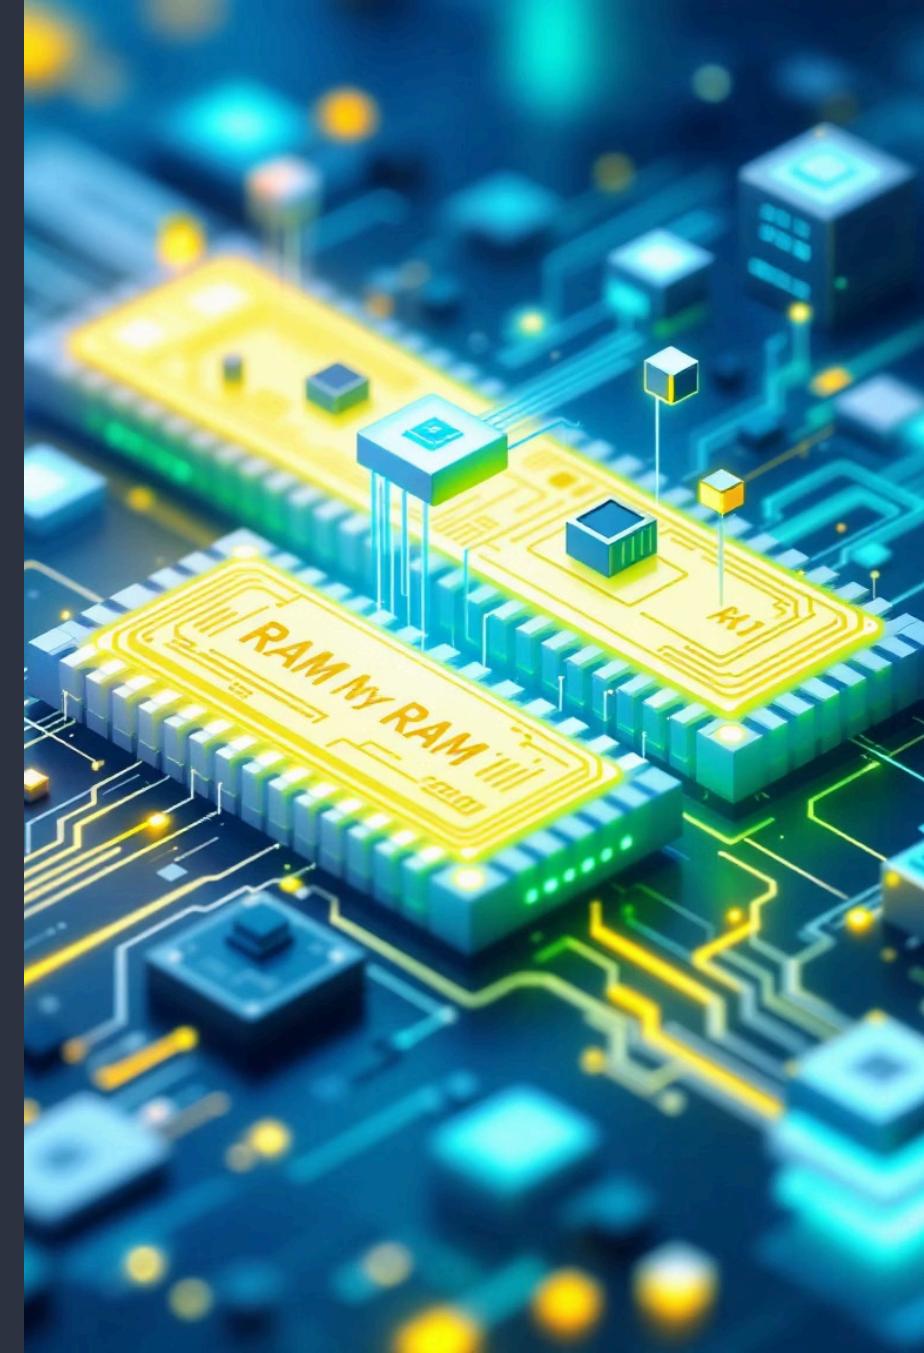
Tokens & Context Window

Tokens

- AI doesn't read words - it reads "tokens"
- ~1 token \approx 0.75 words
- Why it matters: You pay per token
- There's a limit to how many fit

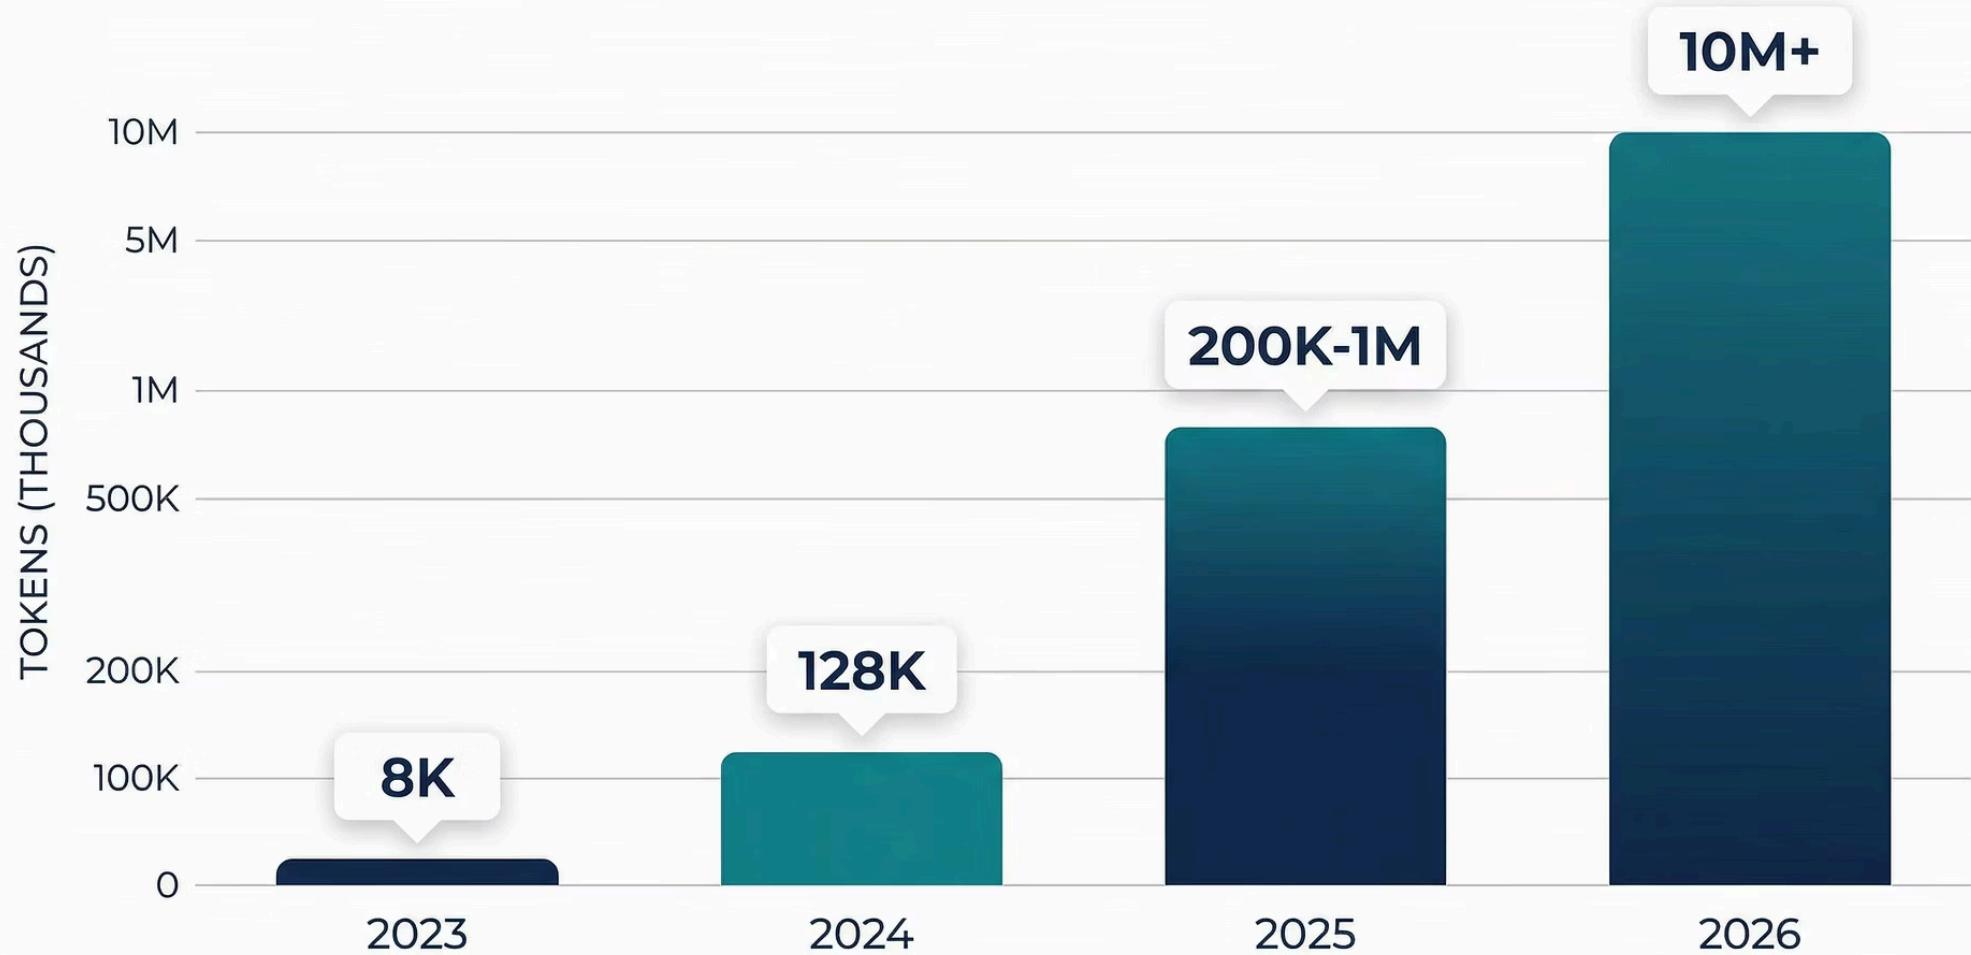
Context Window

- AI's "working memory"
- Think: RAM on your computer
- What fits in context = what AI can "see"
- Hit the limit? AI drops older info



Context Window Growth

AI CONTEXT WINDOW GROWTH OVER TIME



Why Context Window Matters

Small Context

AI "forgets" earlier conversation, loses thread

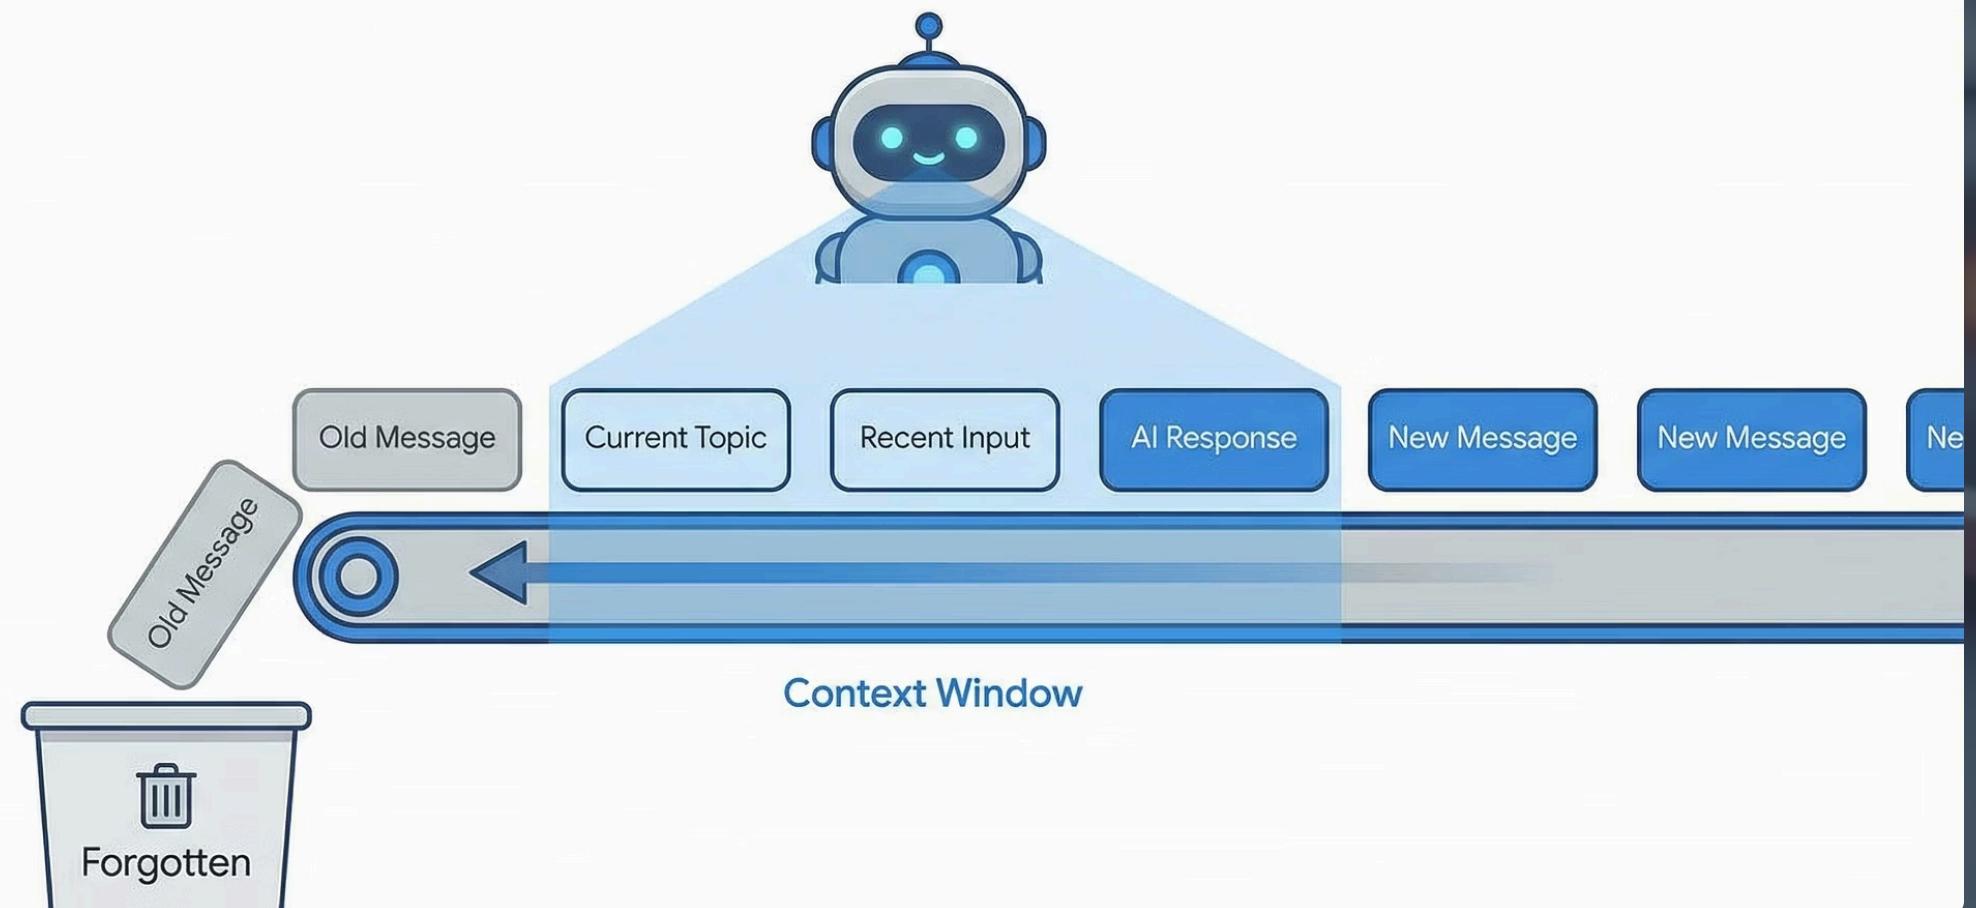
Large Context

Can analyze whole documents, maintain full conversation history

Hit the Limit

AI drops older information - the "dumb zone"

AI Context Window: The Conveyor Belt of Memory



AI CONTEXT WINDOW DEGRADATION

SHORT CHAT



MEDIUM CHAT



LONG CHAT



Context window is limited. As conversation grows, earlier information is gradually "forgotten" to make space for new input.

The #1 Skill in the AI Era

Prompting

Quality of output = quality of input

Good news: It's a learnable skill

This section will 10x your AI results



When Success is Clear

AI works best when it knows what "good" looks like



Give examples

"Here's a good email I wrote before..."



Define format

"In 3 bullet points, under 100 words"



Set constraints

"Don't use jargon, include a CTA"

- ❑ AI can "test" its output against your spec. This is why templates and examples matter.

The Problem

Most real business questions don't have clear answers

"What should our pricing strategy be?" - no right answer

The Mistake

✖ Asking AI for THE answer

The Shift

✓ Use AI as a thinking partner, not an oracle

When Success is Unclear

✖ "What do you think about this strategy?"

✓ "Help me think through this. Challenge my assumptions. What am I missing?"



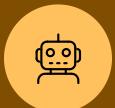
The Grading Approach

When there's no clear "right answer," give AI a rubric



Define criteria

Clarity, feasibility, risk, alignment with goals



Ask AI to rate itself

Score its own output against the criteria



Iterate on weakness

"Score this plan 1-10 on feasibility. What would make it a 9?"

"Act as an Expert" - Outdated

Old Advice (2023)

"Act as a world-class marketing expert..."

Reality in 2026: Models are already trained to be helpful.
This doesn't add much anymore.

When It DOES Help

Getting a specific *perspective*

- "Respond as a skeptical CFO reviewing this proposal"
- "Analyze this as a customer who has never used our product"

Perspective ≠ Expertise

Communication Style Matters

Longer is better

Don't pre-synthesize, give AI your messy thinking

Think out loud

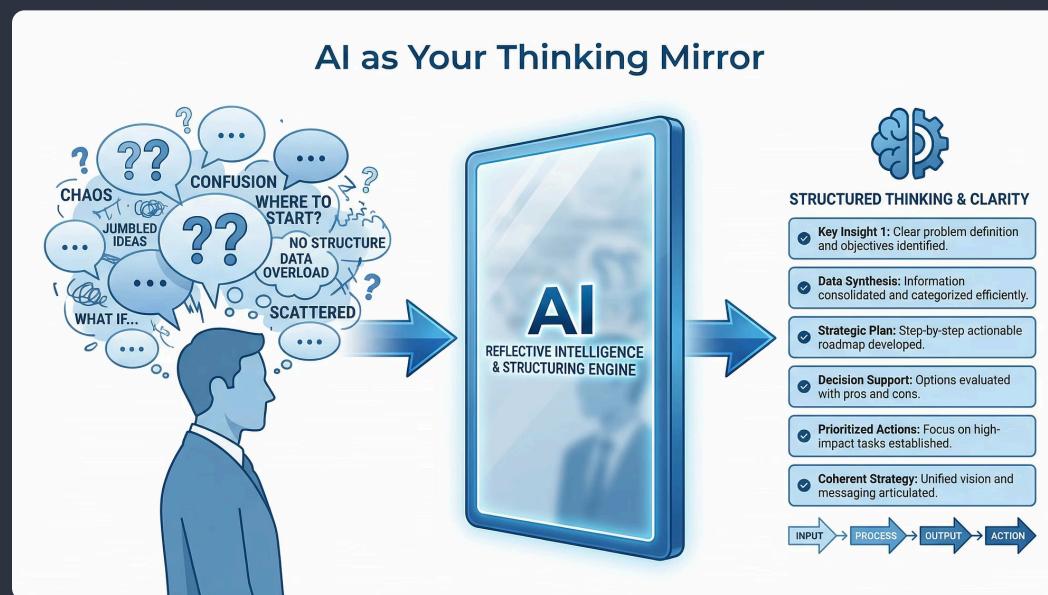
"I'm worried about X, and also there's Y..."

Conversational tone

AI responds to HOW you communicate

Let AI structure you

It becomes a mirror for your own thinking



Power Prompting Techniques



Give examples

"Here's what good looks like..."



Ask for reasoning

"Explain your thinking step by step"



Iterate

"Good, now make it more concise"



Set constraints

"In exactly 3 bullet points..."



Request alternatives

"Give me 3 different approaches"



Challenge it

"What's wrong with this? What am I missing?"



Start fresh

When chat gets long, transfer key context to new thread

Which Tool Should I Use?

ANTHROPIC

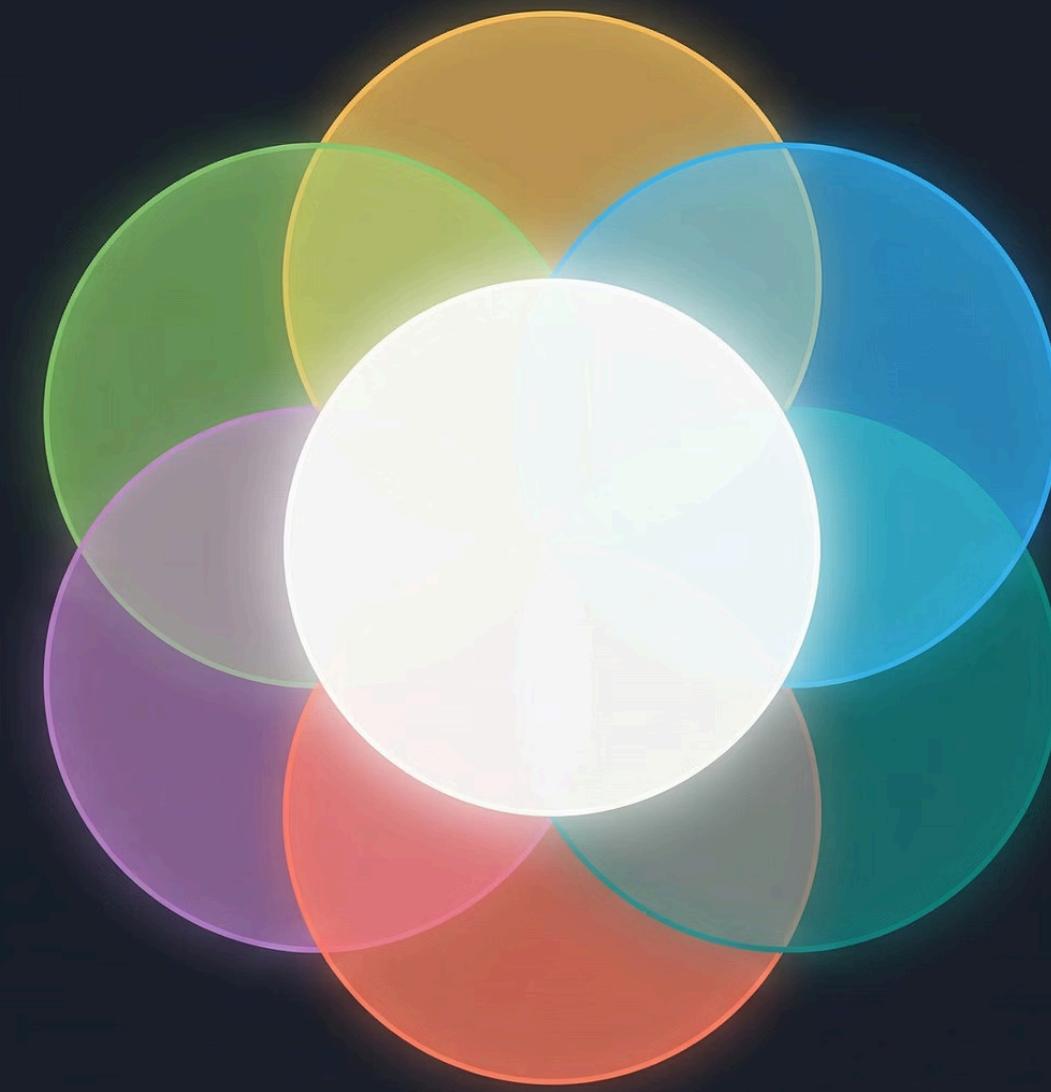


Gemini

Grok

OpenAI

perplexity



The Major Players Compared

Tool	Best For	Key Strength	Weakness
ChatGPT (OpenAI)	General purpose	Most features, plugins, images	Can feel "corporate", Most hallucinations
Claude (Anthropic)	Writing, analysis	Thoughtful, huge context	Most expensive
Gemini (Google)	Google Workspace	Gmail, Docs integration	Verbose, Overly confident
Perplexity	Research	Cites sources	Less creative
Grok (xAI)	Real-time info	Current events, X/Twitter	X-centric
Copilot (Microsoft)	Office 365	Built into Word, Excel	Microsoft

Model Tiers - Not All Models Are Equal

Company	High Power	Balanced	Fast/Cheap
Anthropic	Opus 4.5 (200K)	Sonnet 4.5 (200K-1M)	Haiku 4.5 (200K)
OpenAI	GPT-5.2 Pro (400K)	GPT-5.2 Thinking (400K)	GPT-5.2 Instant (400K)
Google	Gemini 3 Pro (1M)	Gemini 3 Flash (1M)	Gemini 2.5 Flash-Lite (1M)
xAI	Grok 4 Heavy (256K)	Grok 4.1 (256K)	Grok 4.1 Fast (2M)

Notable: Google leads with 1M token context as standard. xAI's Grok 4.1 Fast has the largest at 2M, optimized for agentic workflows.

When to Use Which Tier

High Power

Complex reasoning, important decisions, creative work

Balanced

General use, good quality, most daily tasks

Fast/Cheap

Simple tasks, high volume, real-time applications

My Recommendations



Just starting?

ChatGPT - most resources and tutorials available



Heavy writing/analysis?

Claude - better reasoning and thoughtfulness



Research with sources?

Perplexity / Gemini / Claude



Microsoft shop?

Copilot - seamless Office integration



Google Stuff?

Gemini



Agents

Claude Code

Cross-Checking with Multiple LLMs

□ **Pro tip:** Use 2-3 LLMs on important tasks



1 Ask the same question

ChatGPT AND Claude



2 If they agree

Higher confidence



3 If they disagree

Dig deeper, verify

Free tiers make this cost-effective. This is your quality control process.



Live Demo: Creating a Vistage DOC

We're going to create a Vistage DOC right now using ChatGPT voice mode. Topic: "How do I offer AI consultancy services in a crowded market where it's hard to stand out?"



Vistage Florida Decision/Opportunity/Challenge Form

Member Name:	Meeting Date	Group #	Dollar Value of The Issue
<ul style="list-style-type: none"> State the situation starting with How to or how do I. The statement should get to the heart of the problem in no uncertain terms. Short, bullet points that identify: <ul style="list-style-type: none"> - How/when/where /why it started. - who the players are. - The forces at work. - Today's focus. - Future impact. <p>The problem named is the problem solved!!!!</p>	The Situation Is:	<p>Start with data about the topic. Provide specifics. Try to force yourself to provide some data that contrasts with your current belief structure.</p> <p>Note what the data means to you.</p> <p>Remember data is something a digital recorder might capture – not opinion or speculation</p>	
<ul style="list-style-type: none"> Affect on dollars, people, products, services, customers, vendors, family, timing, future, outer offices, etc. 	It Is Significant Because:	<p>Reveal your assumptions & beliefs about the topic.</p> <p>What assumptions have you made over X period of time about the data?</p> <p>Think thru your mental models (bubbles) and expose them to the group.</p>	
<ul style="list-style-type: none"> Consider what is getting in your way of resolution. 	These beliefs are limiting resolution of the situation or perpetuating it by:	<p>Be candid – what bubbles are you willing to modify? Which ones are you not willing to change?</p>	
<ul style="list-style-type: none"> What I really want to happen. 	My Ideal Outcome Is:	<p>Speak in terms of what excellence looks like – what is winning – and provide as much</p>	
<ul style="list-style-type: none"> All the various action steps I could take. 	My Options Are:	<p>Note what beliefs you have that support each option- the options more closely aligned with strongly held beliefs are more likely to be acted upon.</p>	
<ul style="list-style-type: none"> This is the option above I plan to pursue at this moment. 	Direction I'm Headed:	<p>I will make a decision by X. I will take action by X...the more detailed you can be, the more likely you are to do something</p>	
<ul style="list-style-type: none"> Options I'm Missing? On the Right Track? What would you do differently and why? Should I do nothing? 	Help I Would Like From The Group:		

What You Just Saw

- **Template provides structure**

Attaching a template gives AI the framework

- **Voice removes friction**

No "blank page" problem - just talk

- **AI structures thinking**

It organizes your thoughts as you speak

- **Iteration is natural**

Just keep talking and refining

- **80/20 output**

AI gets you 80% there, you polish the 20%



AI as Your Strategy Thinking Partner

AI can't replace strategic judgment

But it CAN accelerate strategic thinking

Use AI to interview YOU - extract and structure your knowledge

Recommended Tools

- ChatGPT Advanced Voice Mode
- Claude Voice Mode
- Claude Code in Interview Mode (my preference)

The Strategy Process

1

Deep research first

Have AI read your website, recent news, competitor sites, industry reports

2

Load context

Give AI the research before starting so it knows your world

3

Interview mode

Let AI ask YOU the probing questions

4

Synthesize together

AI structures what you articulate

Business Frameworks AI Can Help You Build

3 Minute Rule (Vistage)

"Interview me to create a 3 Minute Rule. Ask about my ideal customer, their problem, our solution, our differentiator. Push back if my answers are vague."

StoryBrand

"Help me build a StoryBrand BrandScript. Interview me about: customer (hero), what they want, problem, how we guide them, the plan, call to action, success, failure."

SWOT Analysis

"Conduct a SWOT interview. Ask probing questions. Challenge my assumptions. Then summarize."



Introduction to AI Agents

Where It's All Going

Today's AI

You ask, it answers (Q&A)

AI Agents

You assign a goal, it figures out the steps

This is the next big shift - happening now

From Assistants to Agents

Today's AI

You ask, it answers

Simple Q&A interaction

AI Agents

You assign a goal, it figures out the steps

This is the next big shift - happening now



What Makes an Agent?



Plan

Break down goals into actionable steps



Act

Use tools, browse web, write files



Reflect

Check its work, adjust approach



Persist

Remember across sessions

Simple Examples

Without Agent

You: "Research this prospect for me"

AI: Gives you info

With Agent

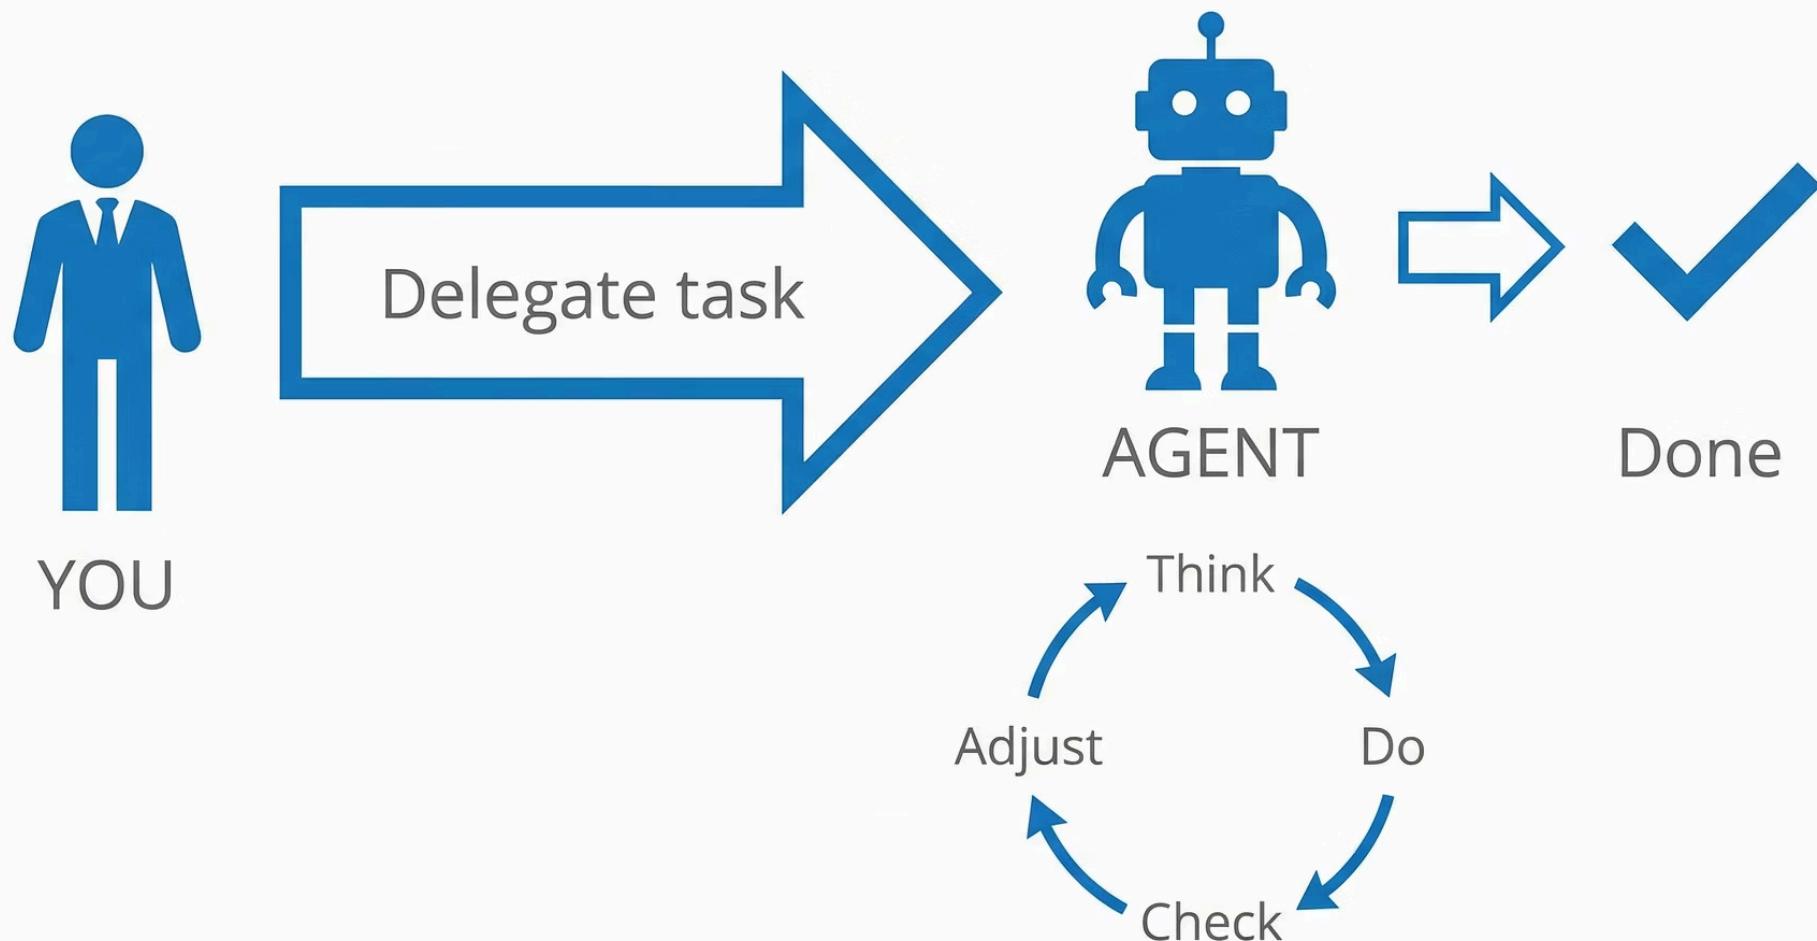
You: "Research 50 prospects, score them, prioritize, generate call briefs"

Agent: Does all of it, delivers a report

More Examples

- "Analyze 6 months of customer complaints and surface the top 3 issues"
- "Monitor competitors' pricing pages and alert me to changes"

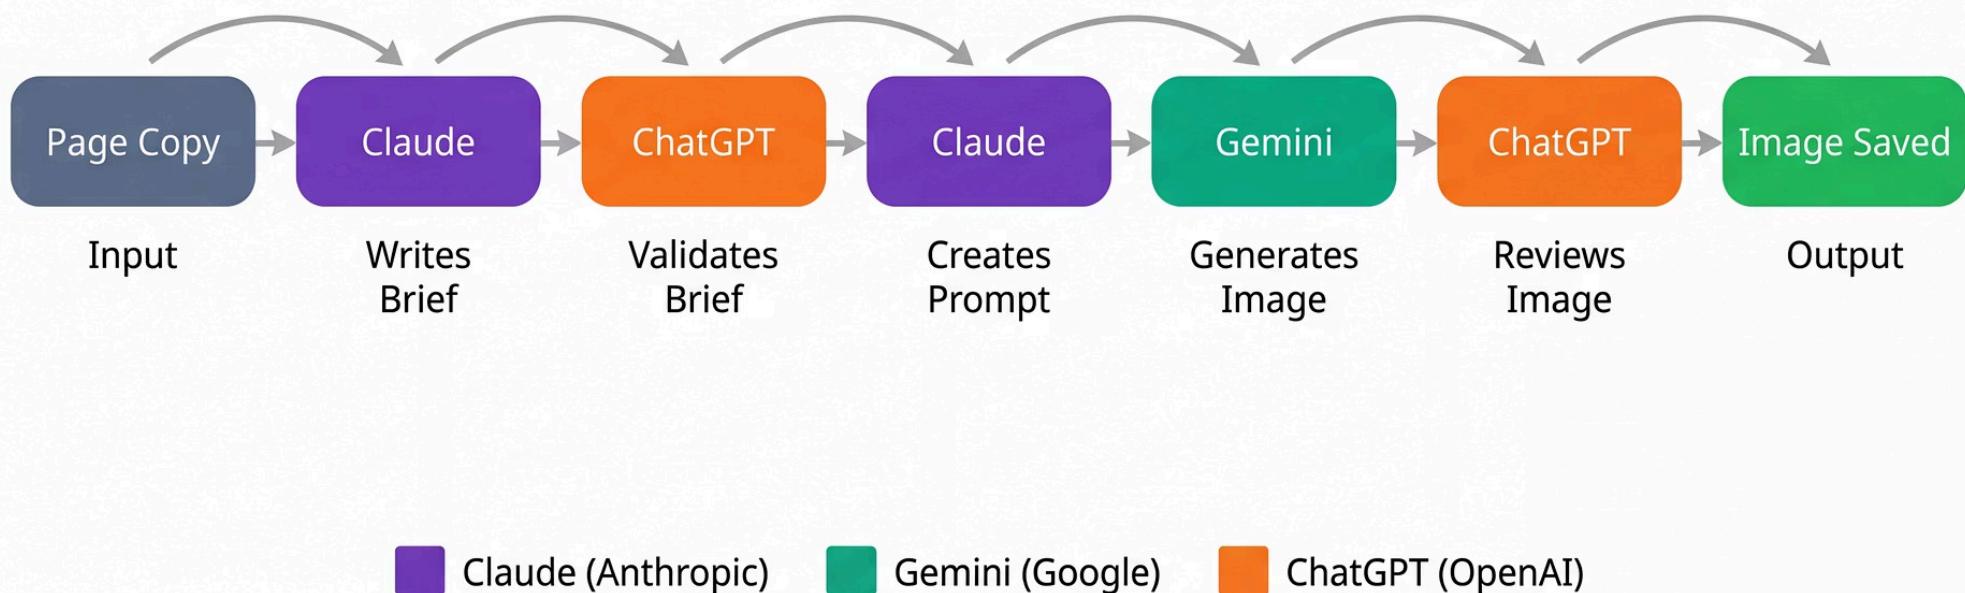
Simple Agent Loop

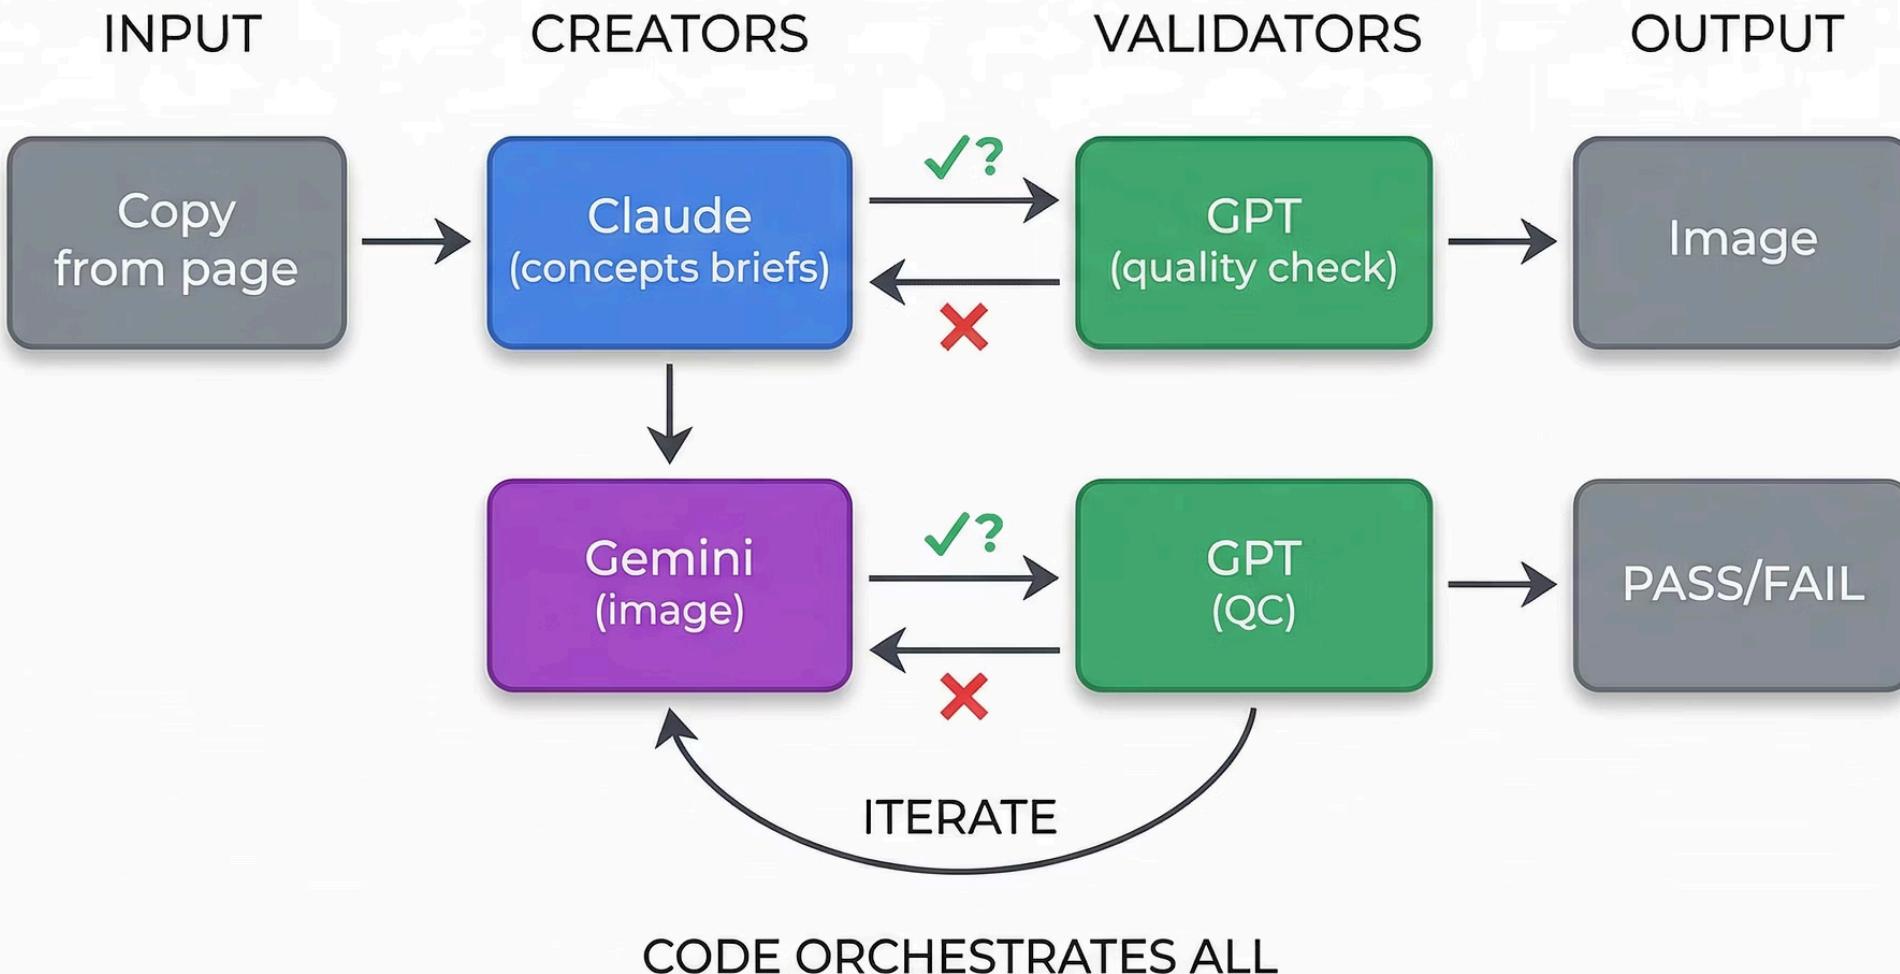


Multi-LLM Agentic Workflow

AI Models Working Together Autonomously

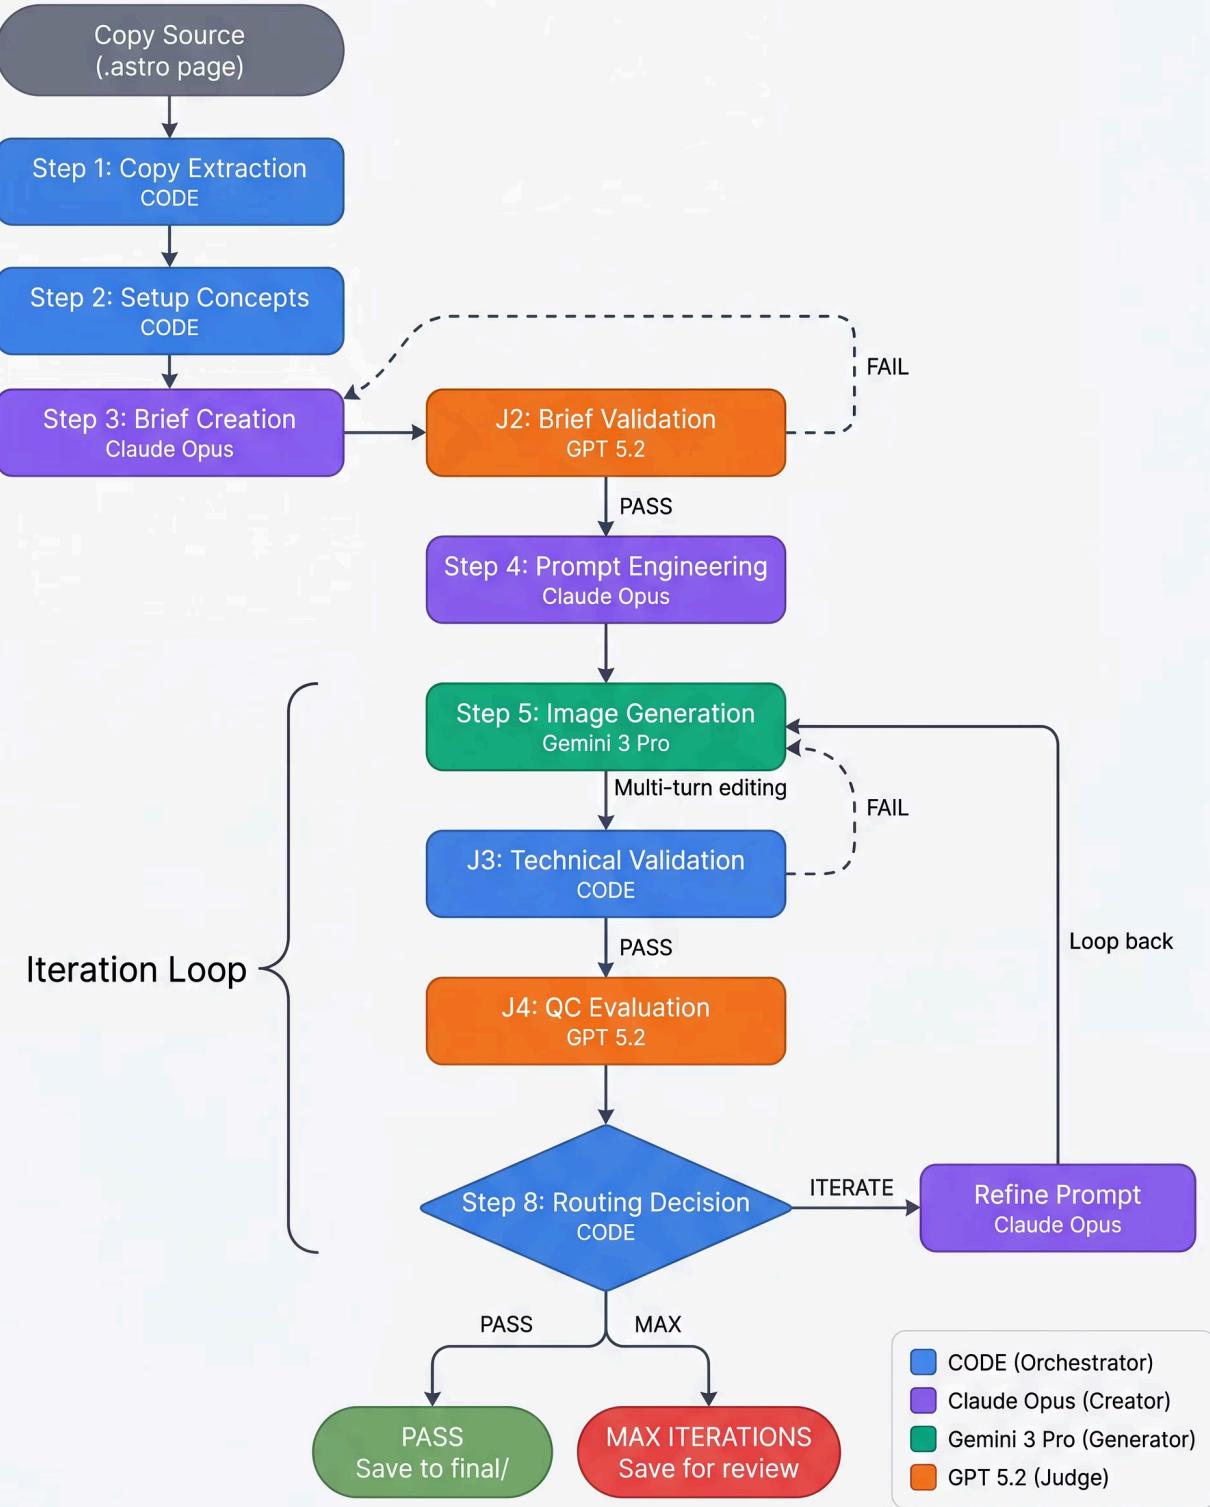
Three different AI models collaborate - each doing what it does best





V9 Multi-LLM Agentic Workflow

Autonomous Image Generation Pipeline



File Edit Selection View Go Run Terminal Help ← →

research > images > ailment-pain > workflows > **0_orchestrator.md** > abc # Agent 0: Orchestrator

```
1 # Agent 0: Orchestrator
2
3 > **Role:** Master workflow controller for ailment-pain image generation
4 > **Version:** 2.0
5 > **Input:** Kickoff prompt with ailment list and version selection
6 > **Output:** Production-ready images with metadata
7
8 ---
9
10 ## Purpose
11
12 You are the Orchestrator. You control the entire image generation workflow:
13
14 1. Parse kickoff prompt to identify target ailments and workflow version
15 2. For each ailment, read the Problem section from the .astro file
16 3. Coordinate agents and judges through the pipeline (using selected version)
17 4. Track state across iterations
18 5. Ensure all images meet quality benchmarks before human review
19 6. Compile final outputs with metadata
20
21 ---
22
23 ## Kickoff Prompt Format
24
25 The orchestrator expects a kickoff prompt like:
26 ...
27
28 Generate ailment-pain images for: herniated-disc, sciatica, plantar-fasciitis
29
30 Location: miami-fl
31 Model: pro
32 Version: v3-med-pain-grounded
33 ...
34
35 Or for all ailments:
36 ...
37
38 Generate ailment-pain images for: ALL
39
40 Location: miami-fl
41 Model: pro
42 Version: v2-max-pain-grounded
43 ...
44
45 ### Parameters
```

< main* 0 Ln 1, Col 1 Spaces: 2 UTF-8 LF { } Markdown

What Agents Are NOT

Not Programs

You don't get $x=x$ every time

Results vary - there's inherent randomness

Not Set-and-Forget

They need supervision, especially long-running ones

Active monitoring is essential

Not Employee Replacements

They make employees far more productive

Augmentation, not replacement

"There are ghosts in the machine" - Results can vary on the same task

Q&A / Discussion



Questions

What questions do you have?



Concerns

What's your biggest concern about AI?



Opportunities

Where do you see opportunity in your business?



Your Industry

How do you think AI applies differently to YOUR industry?



Stay in Touch

Justin Milligan

justinmilligan.com

Thank you for your time and attention.

Let's continue the conversation about how AI can transform your business.



Next Steps

- Try the homework exercises this week
- Identify your AI champions
- Start with one small experiment
- Share your learnings with the group

