




TIER 1 - FOUNDATIONS * V1.0 -- MAY 2026

FREE TOOLS, REAL DIFFERENCES

A practical, opinionated comparison of the four major consumer AI tools. What each one's actually best at, when to switch, and the privacy difference that most reviews skip.



BY

Alex Jahn / Agent Logic

v1.0 -- May 2026

Anyone deciding between ChatGPT, Claude, Gemini, and Copilot -- and tired of the breathless "best AI" articles that compare them on the wrong things

15-20 minutes

Free. Forever.

EDITION

AUDIENCE

READ TIME

COST

Prepared by Agent Logic / alexanderjahn79@icloud.com / theaiguwywi.com

CONTENTS

What's in here

- 1 Stop reading the rankings 3**
Every six months a new "best AI" article ranks the major tools on benchmark scores. Every six months the rankings flip. The tool at #1...
- 2 The four tools in plain language 4**
Quick orientation. The space moves; the personalities are stable.
- 3 The four-question test for any specific task 5**
You have a thing to do. Which tool? Run these four questions in order. Stop at the first clear answer.
- 4 The switching pattern -- two-tool-minimum 6**
Most working AI users I know run two tools, sometimes three. Not because they're paying for everything -- but because the free tiers stack.
- 5 The privacy difference (the part most reviews skip) 7**
This matters more than benchmarks. The four companies have different stances on what they do with your data. As of mid-2026:
- 6 What to skip -- features that matter less than the marketing says 8**
Each tool has a list of "amazing features" you'll see in the marketing. Most of them aren't where the daily value comes from.
- 7 When NOT to use any of these 9**
Three categories where the major consumer chat AIs are the wrong tool:
- 8 Where to go from here 9**
You have the framework for picking. The Tier 1 expansion has three more modules ahead:

SECTION 1

Stop reading the rankings

"Best AI" articles age like milk

Every six months a new "best AI" article ranks the major tools on benchmark scores. Every six months the rankings flip. The tool at #1 last spring is at #3 this spring. The one nobody mentioned a year ago is now everyone's favorite.

The benchmark race isn't useless. It's just not what most users actually need. What you actually need is: when you're sitting at a desk with a question, which tool do you open? You probably already pay for one. You don't want to read 4,000 words about MMLU scores. You want to know -- for *your* questions -- which one wins.

This module is that. It's the framework for picking, not a snapshot of any specific ranking. Snapshots rot in six months. The framework doesn't.

What you'll have by page 13

By the end of this primer:

- A clear-eye Claude, Gemini, Copilot -- on the things that actually matter day-to-day.
- The four-ques
- A switching
- The privacy d
- A what-to-s

The "best AI" question is the wrong question. The right question is "best for what you're doing right now," and the answer changes by the task -- not by the year.

SECTION 2

The four tools in plain language

Quick orientation. The space moves; the personalities are stable.

ChatGPT (OpenAI) -- the original, still the most polished consumer experience. Best general-purpose conversation. Strong on quick answers, image generation, voice mode, and the broadest plugin / app ecosystem. Free tier is real and useful. Paid tier (\$20/mo) unlocks longer conversations, better models, and image features. ChatGPT is the default for most people, and not by accident.

Claude (Anthropic) -- the writer's tool. Best at long-form writing, reasoning through complex problems, working with documents, and not sounding like a sales-y consultant. Free tier is real but more limited. Paid tier (\$20/mo) unlocks the full reasoning, longer context, and a memory feature. If you write for a living or you do work that needs careful thinking, Claude is the move.

Gemini (Google) -- the integrated tool. Best at anything that touches Google's ecosystem: Gmail, Calendar, Docs, Drive, Maps. Free tier is generous. Paid tier (\$20/mo, called Gemini Advanced) is bundled with extra Google storage and adds longer context. If you live in Google's products, Gemini is the lowest-friction path to AI integration with the rest of your day.

Copilot (Microsoft) -- the work tool. Best at anything that touches Microsoft 365: Word, Excel, Outlook, Teams, PowerPoint. Free tier is real (Bing Chat with Copilot, free in Edge). Paid tier (\$30/mo on top of Microsoft 365) unlocks Copilot inside the Office apps. If your day is Excel and Outlook, Copilot's office integration is genuinely the highest-leverage purchase. If your day isn't, Copilot is the weakest of the four for general use.

The one-liner per tool:

- **ChatGPT** -- the best general-purpose AI assistant.
- **Claude** -- the best AI for actual writing and reasoning.
- **Gemini** -- the best AI if you live in Google.
- **Copilot** -- the best AI if you live in Microsoft Office.

If your situation matches one of those one-liners cleanly, you have your default. The rest of this module is for the people whose situation doesn't.

SECTION 3

The four-question test for any specific task

You have a thing to do. Which tool? Run these four questions in order. Stop at the first clear answer.

The four-question test:

- 1. Does this task touch a specific ecosystem?** Gmail, Calendar, Drive -> Gemini. Office apps -> Copilot. Otherwise, keep going.
- 2. Is this writing or reasoning that has to be precise?** Drafting a long document, working through a hard decision, summarizing a complex thread, editing real prose -> Claude.
- 3. Is this a quick question, an image task, or a casual conversation?** ChatGPT (free tier covers most of these).
- 4. Is this experimental or for fun?** Try whichever one you haven't used in a while. The free tiers are good enough for play.

That's it. Most tasks resolve at question 1, 2, or 3. The fourth question is a reminder that all four tools have improved enough that your old assumptions are probably stale.

Worked tests

Task: write a follow-up email to a client about scheduling a project.

- Question 1: not Gmail-shaped specifically (you're writing the content, not navigating Gmail). Skip.
- Question 2: yes -- actual writing where the words matter. -> **Claude.**

Task: summarize a long thread of emails before a meeting.

- Question 1: yes -- Gmail. -> **Gemini.** It pasting it).

Task: clean up an Excel sheet that has a column of half-formatted dates.

- Question 1: yes -- Office. -> **Copilot.** It

Task: figure out why your truck's check-engine light came on.

- Question 1: not ecosystem-shaped. Skip.
- Question 2: not writing-shaped. Skip.
- Question 3: quick question. -> is solid.

ChatGPT

Task: brainstorm names for a new business.

- All four tools are fine. Question 4: try the one you haven't used in a while. Variety surfaces ideas your default tool wouldn't.

SECTION 4

The switching pattern -- two-tool-minimum

Most working AI users I know run two tools, sometimes three. Not because they're paying for everything -- but because the free tiers stack.

The cheap stack that covers most needs

Pay for one, free-tier two. Pick one paid tier (\$20/mo) and use the free tiers of the others as needed.

For most non-Office-heavy users, the highest-ROI paid tier is differentiator, the writing quality is real, and Claude doesn't have a strong free competitor at the writing-quality tier. Use ChatGPT free for quick questions and image work; use Gemini free for Google integration when needed; pay for Claude.

Claude P

For Office-heavy users, already use eight hours a day. Use ChatGPT free or Gemini free for the rest.

Copilot in

The setup, once

- Install the app or save the bookmark for all four -- this is free. ChatGPT, Claude, Gemini (or just gemini.google.com), and Copilot (built into Edge or Microsoft 365 if you have it).
- Decide your default per the one-liners in section 2 and the four-question test.
- Pay for one. Don't pay for two unless you've already proven you switch heavily.

That's the cheap stack. Total cost: \$20-\$30/mo. Total coverage: roughly 95% of consumer AI use cases.

Pay for one tool, free-tier the others. The "should I subscribe to two" question almost always resolves to "no, switch instead." Switching is free. Subscribing twice is \$20/mo of waste.

SECTION 5

The privacy difference (the part most reviews s

This matters more than benchmarks. The four companies have different stances on what they do with your data. As of mid-2026:

OpenAI (ChatGPT) -- by default, conversations on the free tier may be used to train future models. Plus and Pro tiers have a setting where you can turn this off (Settings -> Data Controls).

Enterprise/Team tiers don't train on your data by default. Practical implication: if you're on the free tier and you don't want your conversations used as training data, either switch to a paid tier with the toggle off, or don't share sensitive content.

Anthropic (Claude) -- by default, Claude does NOT train on consumer conversations (free or Pro). This is a real differentiator. They use conversations only for safety review and to flag policy violations, not for model training. Practical implication: among the four, Claude has the cleanest default privacy posture for individual users.

Google (Gemini) -- by default, conversations may be reviewed by humans for quality and may be used to improve services. There's a setting to turn off conversation history (Activity controls), but turning it off also weakens the assistant's usefulness over time. Practical implication: assume Google sees the data, treat accordingly. Gemini is the strongest tool for ecosystem integration but the weakest for sensitive personal content.

Microsoft (Copilot) -- varies by tier. Consumer Copilot has commercial-data-protection on by default for signed-in users, which is genuinely better than the consumer norm. Microsoft 365 Copilot (the work version) has enterprise-grade controls. Practical implication: Copilot is reasonable for personal use, especially compared to free tiers of the others.

The simple rule

If you're going to share sensitive information with an AI, default to paid tier of the others with training-data toggles confirmed off.

Claude (b

If you're going to share a client's data, an employee's data, or any third-party private information -- none of these consumer tools is the right tool. That's a business-tier conversation: Microsoft 365 Copilot, ChatGPT Enterprise, Claude for Work. The free consumer tier is for your own questions.

The privacy red line:

- Personal preferences, opinions, half-formed ideas -> any tool, fine.
- Your own non-sensitive professional content -> any tool, fine.
- Sensitive personal content (health, family, money) -> Claude or paid tier with training off.
- Other people's private information -> not a consumer-tier conversation. Use a business tier or don't use AI for that task.

SECTION 6

What to skip -- features that matter less than the

Each tool has a list of "amazing features" you'll see in the marketing. Most of them aren't where the daily value comes from.

Image generation -- fun, occasionally useful, not a daily driver. ChatGPT (DALL-E) and Gemini both do it. Claude doesn't. Don't pick a tool regularly. *for image*

"Agents" / "deep research" / "deep think" buttons -- every tool now has a "think harder for longer" mode. Useful sometimes; takes 5-15 minutes per query. Use sparingly when you need depth; don't make it your default.

Custom GPTs / Gems / Custom Claude Projects -- these are pre-configured personas you can build. Genuinely useful if you have a recurring task; ignorable if you don't. Most people overestimate how much they'll use them.

Voice mode -- covered in detail in situational. Don't pick a tool *AI On The only for vo*

Code generation -- the major tools all do this competently. If you're a working developer, the gold standard is a developer-targeted IDE integration like Cursor or GitHub Copilot -- not the consumer chat tools.

Mobile vs desktop apps -- all four have both. Mobile apps are now solid for all four. Don't pick a tool because of "the app is good" -- they all are.

The honest list of features that *do matter:* privacy posture, and free-tier generosity. Everything else is decoration.

SECTION 7

When NOT to use any of these

Three categories where the major consumer chat AIs are the wrong tool:

- **High-stakes professional work.** Medical, legal, tax. The tools will give you confident-sounding answers; they're sometimes wrong, often hedged into uselessness. For anything that has consequences, talk to a real professional.
- **Workflows where you'd benefit from automation, not chat.** If you're typing the same prompt every Tuesday, you don't need a better chat tool -- you need a script or a small custom integration. That's Tier 2 / Tier 3 territory.
- **Anything you'd be embarrassed to have logged.** Even with good privacy postures, every conversation is logged on some server somewhere. Don't put anything in a chat window you wouldn't put in an email.

Within those limits, the four tools cover essentially all consumer AI needs. Picking the right one for a task matters more than picking the right one overall.

4

Four questions, four tools.

Run the four-question test on any task and you'll know which tool to open. Pay for one, free-tier the rest.
Switch by task, not by loyalty.

SECTION 8

Where to go from here

You have the framework for picking. The Tier 1 expansion has three more modules ahead:

- **AI for big decisions** -- using AI to compare cars, houses, contracts, services. Where AI lies (specs, prices) and the moves that catch it.

After that: kids, walls of text, and Tier 1 is done at 12 modules.

Get the next module the day it drops: theaiguywi.com/training

One email per release. No drip. No spam. Opt out anytime.

If you want this same comparison framework installed across a team -- the right tool taught for the right job, the privacy red lines drawn before someone shares the wrong thing, the cheap-stack pattern adopted shop-wide -- that's the consulting offer. We do it the way I run it in my own carpentry business.

Reach out: alexanderjahn79@icloud.com

A short call. Honest scope. We figure out together if it's a fit.

Closing -- the lock-in line

The "best AI" question wastes time because the rankings change every six months and the answer depends on the task anyway. The four-question test outlasts the rankings. Run it on any task. The right tool will surface in under ten seconds.

The pay-for-one, free-tier-the-rest pattern is the cheapest way to cover 95% of needs. Don't subscribe to two. Switch instead.

You have the comparison. Three more modules in this batch.

Agent Logic --

Fond du Lac, WI. This is module 9 of 12 in Tier 1 (Personal).

© 2026 Agent Logic. Share freely.

theaiguyw