

AI Marketing Tools 2.0

What changed, what one person built, and what the buyer's AI says about you

AHFA Marketing/PR Conference · August 26, 2026

Matt Ritter · Nido R. Qubein School of Communication · High Point University



The Gap Between Reality and Aspiration

THE REALITY

40 hrs/wk

Same hours. Same headcount.
Increasing expectations.

AI

THE ASPIRATION

- Follow up with every email promptly
- Prep thoroughly for every meeting
- Personalize every student touchpoint
- Document what actually works. Fix what is broken.
- Optimize my time
- Offload labor-intensive tasks

BEFORE AI AFTER AI

IDEA



EXECUTE



IDEA



EXECUTE



Where We Left Off

Last year, in thirty seconds.

AI AS INTERN

Eager, fast, solid first drafts,
sometimes confidently wrong.
Your expertise is the quality
control layer.

THE TOOLS

Claude, ChatGPT, Copilot,
Gemini. Learn to prompt one
well and you can use any of
them.

PROJECTS

Teach it your business once.
Brand voice, product specs,
examples of your work. Then
just ask.

What Changed Since We Last Met

Four shifts in twelve months

1 Memory

All four major tools now remember prior work by default. You stop re-explaining.

2 Agents

Claude Cowork and ChatGPT agents move from answering to executing multi-step tasks and reporting back.

3 Copilot runs Claude

Microsoft Copilot now offers Claude models alongside GPT inside Word, Excel, Outlook, and Teams.

4 Search inside the tools

Every major assistant searches the live web and cites its sources.

The Tools You Can Use Today: My Favorite Interns

MY GO-TO

Claude

Anthropic

Strong reasoning, nuanced writing, Projects for context. Cowork automates workflows.

\$20/month

MOST POPULAR

ChatGPT

OpenAI

The original. GPT-5 models, image generation, voice, agents for automated workflow.

\$20/month

OFFICE INTEGRATION

Copilot

Microsoft

Built into Word, Excel, Outlook, and Teams. Powered by GPT and now Claude.

**\$20/individual;
\$30/consumer**

GOOGLE INTEGRATION

Google Gemini

Google

Built into Gmail, Docs, Sheets, and Search. Seamless for Google Workspace users.

\$20/month

Last week, I gave the interns a test.

“My car is filthy and needs to be washed, but the car wash is 50 meters away. It’s so close. Should I walk or drive?”

The Ones Who Told Me to Walk

GPT 5.5 QUICK

“Walk. 😊 ...driving there would save you roughly 30 seconds of walking while requiring you to start the car, shift into gear, and travel a distance that many people could cover with a good sneeze.”

“My recommendation: walk to the car, drive it to the car wash, then feel slightly ridiculous for having considered driving the 50 meters yourself. 🚗💨🚶”

COPILOT QUICK

“Walk. 😊 ...this is one of those rare situations where walking is objectively the more efficient choice. Everybody wins. 🚶🚗💨”

CLAUDE OPUS, VIA COPILOT

“Walk, without a doubt! ...Lace up and walk. Your legs (and the environment) will thank you! 😊”

The Ones Who Drove

CLAUDE OPUS, ASKED DIRECTLY

“Drive. The car has to come with you. Walking there gets you a clean pair of shoes and a car that’s still filthy 50 meters away.”

GPT 5.6 DEEP THINK

“Drive. The point is to get the car to the wash, not just yourself. 😊”

Copilot Deep: “Walking there would leave the filthy car at home. 🚗💧”

HONORABLE MENTION: CLAUDE SONNET, MID-SENTENCE

“Cons: You’ll need to either push/roll the car there somehow (not really an option 😊) or... wait, you do need to drive the car to get it washed, don’t you?”

The quick modes walked. The thinking modes drove. And the same model gave different answers depending on the door I knocked on.

Two Things I Often Hear about AI

“I HIT LIMITS TOO FAST”

Still true. Heavier models and agents burn through allowances faster, and the \$20 tier caps out on a busy morning.

- Match the model to the task
- Save the heavy model for reasoning and drafting
- Expect to hit the ceiling on the day you need it most

“I DON'T WANT COMPANY DATA IN THE CLOUD”

Team and enterprise plans are how most companies solve this. Data stays out of training, admins control access, and Projects can be shared across a department.

If your company has not looked at a team plan, this is a good place to start

What's Real Today

Honest calibration: production-ready versus emerging

PRODUCTION READY

- Research assistants
- Writing and editing partners
- Personal productivity agents

Real, used daily, reliable enough to trust.

STILL EMERGING • 12 TO 24 MONTHS

- Multi-step operational workflows
- Regulated, compliance-heavy domains
- High-stakes autonomous action

Coming, but unevenly.

Where Many of Us Actually Are Right Now

Recently, a senior leader admitted that their tasks, ideas, and reminders all lived on sticky notes.

Stuck to the monitor. Stuck to the desk. Stuck to the wall.

“I think I have a system.”

But things were slipping. Important things.

Once everything is a priority,
nothing is.

Meet Erin

Same problem. Different response.

WHO

Erin

ROLE

Chief Financial Officer

THE PROBLEM

Too many competing demands.

- Dropping things
- No real-time consolidation across email and Teams
- No clear answer to: “What do I focus on first?”

She Started Here

Four sentences, written one Tuesday afternoon

THE STARTING PROMPT

Act as my Chief of Staff.

Each morning, give me a prioritized daily briefing based on my Outlook calendar, my emails, and my Teams messages.

Be direct, concise, and ruthlessly prioritized. No fluff.

I want clarity on what I must do today.

Then She Iterated

Each addition came from noticing a specific failure during the day

WEEK 1

Add:

Ignore auto-generated system emails. I'm getting noise from platforms my team no longer uses.

WEEK 2

Add:

Cross-reference email AND Teams before flagging anything as unresolved. Sometimes I respond on Teams to an email request and the briefing thinks it's still open.

WEEK 3

Add:

Scan my sent items from the past 7 days for commitments I made that haven't been completed. "I'll send that over." "Let me check and get back to you." Those are dropping through.

WEEK 4

Add:

For each meeting today with other people, suggest 2-3 sharp questions I should ask, based on the agenda, recent threads, and what's been pending with these people.

Every one of these started as a sentence she could have written into a notebook. Instead, she wrote it into the prompt.

What It Produces Each Morning

Open items with direct reports, from one Tuesday morning's briefing

Sarah Chen

- Waiting on: Q3 forecast model update (4 days open)
- Waiting on: Vendor contract review (12 days open)
- Owe her: Walk through budget revision at next 1:1
- Watch: Quieter than usual on the new initiative

David Park

- Waiting on: Growth plan for Q4 (asked 9 days ago)
- Waiting on: Director-level hire scorecard
- Good work: Caught the reporting error in the board materials. Flag for his annual review.

Maria Rodriguez

- Due today: Month-end close confirmation
- Owe her: Decision on the consultant engagement
- Upcoming: Performance review Friday

*It doesn't just track tasks. It tracks people, in both directions. *Details generalized for this presentation.*

What It Becomes

Eight weeks later, the same prompt, fully grown

4

sentences in the starting
prompt

8

weeks of iteration

32+

specific instructions in
the current version

1

person built it. No IT, no
consultants.

She didn't sit down and write this. She discovered it, one missed email at a time.

Now the Marketing Version

STARTING PROMPT

Act as my marketing admin. Each morning review my inbox, my Teams messages, and today's trade press. Tell me which retailer emails need a response, which competitor or trade news I should know before my 10 a.m., and what I owe people. Be direct. No fluff.

WEEK 1

Stop flagging every *Furniture Today* newsletter. Only surface stories about our top ten accounts, our named competitors, or High Point Market.

WEEK 2

When a rep forwards a retailer request, track it as an open item until I confirm it shipped, even if I answered the rep on Teams.

WEEK 3

Before any account call, pull the last three emails with that buyer and draft two questions I should ask.

WEEK 4

Every Friday, list every commitment I made this week in sent mail that has no follow-up yet.

You don't need to know your workflow before you start.

You need to delegate it badly, watch what breaks, and pay attention.

That's the documentation. That's the asset.

Your Four Sentences

Thirty seconds. Silently. For your own most annoying recurring task.

Act as my _____.

Each morning, _____.

Be _____.

I want _____.

If you can fill in those blanks, you can start Monday.

Part one: Ways to use the intern

Part two: What does an intern say to your customers about your brand when you're not in the room.

Your Buyer Asks Before They Click

The research is moving inside the answer

THE SEARCH PAGE CHANGED

43%+

of U.S. Google searches now show an AI Overview

15% → 8%

clicks on a traditional result when an AI summary appears

THE HABIT FORMED

9.5B

monthly visits to generative AI platforms worldwide, up 70% year over year

6 in 10

ChatGPT referrals land on a homepage.

THE VISITOR WHO DOES CLICK

2.5x

more likely to visit a brand's site within 7 days after ChatGPT recommends it, versus a non-recommended competitor

+393%

AI-referred traffic to U.S. retail sites, Q1 2026 vs Q1 2025.

Sources: Similarweb, 2026 Generative AI Landscape, July 29, 2026. Similarweb, The Downstream Impact of AI Visibility, June 21, 2026. Pew Research Center, July 22, 2025. Adobe Digital Insights, Apr. 16, 2026.

Same Game, New Referee

SEO

Rank on a list of links.

Google reads keywords, backlinks, and page structure.

GEO

Get named inside the answer.

The model reads facts, sources, consistency, and what other people say about you.

Google's own position, May 15, 2026: optimizing for generative AI search is still SEO. AI Overviews and AI Mode retrieve from the same core Search index. Everything you already do still counts.

Source: Google Search Central, "Optimizing your website for generative AI features on Google Search," published May 15, 2026.

What the Machine Rewards

The first large-scale academic test of what gets content cited in AI answers

RAISED AI CITATION VISIBILITY 30 TO 40 PERCENT

- Add statistics
- Add quotations from credible sources
- Cite sources

Also helped: clear, fluent prose and an authoritative voice.

HURT

Keyword stuffing

8% below baseline

- *A page that repeats "mid-priced upholstery" eleven times reads as low-quality and promotional*
- *Models prefer "eight-way hand-tied frames, built in Maiden, North Carolina, lifetime warranty" once, plainly.*

Source: Aggarwal et al., "GEO: Generative Engine Optimization," ACM KDD 2024. Princeton, IIT Delhi, Georgia Tech, Allen Institute for AI. 10,000 queries.

I Ran My Own Test

Two questions. Three tools.

QUESTION 1

“What mid-priced upholstery brands should an independent furniture store carry?”

QUESTION 2

“What are the best American-made case goods manufacturers?”

Asked the same way, the same day, in ChatGPT, Claude, and Gemini. Then I read who got named, and what each tool was reading.

Who Showed Up in Every Answer

Out of roughly 50 brands named across six answers, only six made all three tools

UPHOLSTERY

2 of 26

Rowe
England

Everyone else appeared in one or two answers.

CASE GOODS

4 of 25

Stickley · Gat Creek
Copeland · Vaughan-Bassett

Everyone else appeared in one or two answers.

The question isn't whether AI recommends you. It's whether you're a consensus name or a coin flip.

Where the Answers Came From

Each tool read a different kind of source. Each one is a door you control.

ChatGPT read you.

Manufacturer websites, almost exclusively. Rowe, Craftmaster, Best, Gat Creek, Stickley, Vaughan-Bassett, all cited from their own pages.

**If your site doesn't say it plainly,
ChatGPT won't say it.**

Gemini read your dealers.

Retailer pages above all: Woodstock Outlet, Green Front, Furniture Fair, Tanger's. Plus Houzz, Apartment Therapy, and social.

**The way your dealers describe you is
feeding Gemini's answer about you.**

Claude read the trade press.

Furniture Today's Domestic Dynamos feature drove most of the case goods answer.

**One good trade story can carry a
brand into the answer.**

What the Consensus Brands Have in Common

A CRISP, QUOTABLE CLAIM

Vaughan-Bassett: “100% made in the USA, no exceptions.” Repeated word for word by all three tools.

Also lifted verbatim: Gat Creek “signed by the builder,” Stickley “over 90% domestic”

SPECIFIC PRODUCT PAGES

Rowe, England, and Craftmaster got cited from pages that state construction, fabric counts, and quick-ship terms in plain language.

Facts and numbers on your own pages are what get lifted into the answer.

OTHER PEOPLE SAYING IT TOO

Dealer pages that name and describe them. A trade press feature. A Reddit thread.

Consistency across your site, your dealers, and the press is what a model reads as authority.

Consistent with Aggarwal et al., ACM KDD 2024 (statistics, quotations, and cited sources raised AI citation visibility 30 to 40 percent).

Room for Growth

Two well-regarded NC upholstery makers were absent from all six answers. The facts were there. The machine could not find them.

WHAT I FOUND

- Homepage copy reads “built with integrity, crafted with heart.” No construction, origin, or warranty on the page a model reads first.
- The facts exist one or two clicks down: eight-way hand-tied, 1,300+ styles, lifetime frame warranty, four NC plants, 670 associates. On a “Buying Facts” page and a Careers page.

THE CONTRAST: VAUGHAN-BASSETT

*“100 percent crafted in the United States.
500+ craftsmen in Galax, Virginia. 99
percent of lumber from within 500 miles.”*

Stated on the About page and on every brand page. Then repeated nearly word for word by Woodstock Outlet, Wood’s Mercantile, and All American Made.

The gap isn’t quality, it’s where the facts sit, how they’re phrased, and whether anyone else repeats them.

Spot-check: templefurniture.com, leeindustries.com, vaughanbassett.com, hickorypark.com, cabothousefurniture.com, woodstockoutlet.com, woodsmercantile.com, allamericanmade.com, August 2026.

A GEO Checklist

1 Audit

Ask four tools how they describe your brand and your top collections. Screenshot answers and sources. Repeat monthly.

2 Keep doing SEO

Google says the systems are the same. May 2026 guide.

3 Standardize product data

Same attributes, same labels, every listing.

4 Add facts and citations

Frame construction, fabric specs, certifications, warranty terms on collection pages. (Princeton, KDD 2024.)

5 Earn editorial placement

Wirecutter, Apartment Therapy, The Strategist, and trade press. (SWPR)

6 Show up where consumers talk

Answer real questions on Reddit and design forums, credibly. (Modern Retail)

7 Don't stuff, don't plant

Keyword stuffing hurts. Planted authority statements are being countered. (Princeton and WSJ)

The Line You Should Not Cross

Some firms plant “brand authority statements” across sites and use superlatives to trigger recommendations. The platforms are pushing back. OpenAI cites in-house indexes, Bing integration, and veracity weighting. Google points to ranking systems built to resist keyword stuffing in Gemini.

If a claim would not survive a call from a Furniture Today reporter, do not publish it for a chatbot either.

Source: Christopher Mims, Wall Street Journal, Jan. 30, 2026.

Two Takeaways

FOR YOU, THIS WEEK

Pick one recurring task.

- Write your four sentences
- Watch where it fails
- Articulate the failures back to the AI
- Iterate

Six weeks: something useful. Six months: the expert in your org.

FOR YOUR BRAND, THIS QUARTER

Run the audit first.

- Ask the tools about your brand
- Fix your product data
- Get named on dealer pages and in trade press

The Real Opportunity

The goal isn't fewer people.
It's more, and better, from the same people.

Better preparation. Clearer communication. Faster first drafts. More time for the work that actually requires you.

Questions?

Matt Ritter · High Point University

mritter@highpoint.edu