2025 | OCTOBER

AI Authentic Insights

Discoveries and dispatches from the AI frontier

LYDON

OCTOBER **2025**

Contents

AI_01: FROM THE EXECUTIVE CREATIVE DIRECTOR

The pernicious promise of progress 3

AI_02: FEATURED PERSPECTIVES

Amplifying AI impact while safeguarding trust 6

By Marc Whitt, guest contributor

AI_03: FEATURED PERSPECTIVES

Rocking AI search with a "Triple S" framework 8

By Chris Phenner, guest contributor

Reports from the field 10

AI_04: DEPLOYMENT STRATEGIES

Prime prompting: Ask (the right way) and you shall receive 12

AI_05: GROWTH HORIZONS

Vibe coding is no silver bullet but can still provide useful ammo 14

AI_06: APPLIED SCIENCE

AI can make moments—creatives make them matter 17

AI_07: SPEED TO DELIVERY

Clearing the clutter to make way for great work 19

A working AI glossary 20

AI_08: HOT TAKES

Why AI will probably never get creativity 21

LYDON

The pernicious promise of progress

B | L

Since the invention of the wheel, human history has been defined by invention. From steam engines and electricity to antibiotics and the printing press, our most revolutionary technologies have advanced society, elevated living standards and empowered generations. The Industrial Revolution, in particular, accelerated this progress—delivering railways, factories, telecommunication and mass production.

Even throughout the early 2000s, most major breakthroughs—like GPS, broadband internet, the iPhone and hybrid vehicles—delivered a net benefit to society, despite their imperfections.

But something has changed.

In the past two decades, a troubling pattern has emerged: a growing number of new technologies—especially those rooted in digital ecosystems and artificial intelligence—seem to be doing more harm than good. Rather than solving meaningful problems, they "fix" problems we never had, automate tasks that degrade human agency and concentrate power in fewer hands. This isn't disruption as we once knew it—it's erosion of jobs, trust, identity, attention and even truth.

The rise of technologism

At the heart of this shift is a belief system often called "technologism," or "techno-centrism"—the idea that every problem can (and should) be solved with new technology. Under this worldview, human wisdom, cultural tradition and systemic reform take a backseat to apps, algorithms and automation. Problems like loneliness, inequality, climate change and disinformation aren't addressed with structural solutions, but with product launches and platform releases.

This isn't just misguided—it's dangerous. When we confuse novelty with progress and complexity with superiority, we turn tools meant to liberate us into instruments of control, distraction and decay.

Generative AI: Innovation in overdrive

Take generative AI. Touted as one of the greatest breakthroughs of the 21st century, it's capable of fabricating human-like text, images, audio and code with astonishing plausibility. Yet, rather than ushering in a new creative era, it's flooding the internet with junk content, eroding artistic value, undermining trust in what we see and read, and contributing to a staggering increase in fraud and misinformation. Worse still, it consumes massive amounts of energy, exacerbating an environmental crisis that technologists claim they're trying to solve.

Business leaders have long warned of a widening skills gap in developed economies. Jamie Dimon of JPMorgan Chase believes the solution lies in education, not more technology[1]—and I agree. Let's be honest: how many of us really believe tools that devalue human ingenuity are going to enhance human skills?

Recent research from Apple, "The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity," shows that large AI models still struggle with complex reasoning. In other words, they may sound smart but they actually understand very little. So, before we start relying too heavily on LLMs and LRMs, we should scrutinize the trade-offs and start asking some tough questions: Are we really better off with AI-generated content farms, scam-calling voice clones or deepfakes that could destabilize democracies? Or are we simply too dazzled by what the technology can do that we're ignoring what it should do?

According to the 2024 Edelman Trust Barometer, global trust in institutions—including government, media and even businesses—is in long-term decline. In the U.S., trust in media sits below 40 percent. Just 47 percent of people globally say they trust that innovation is being handled well. In other words, people are losing faith that technology, and those who develop it, are working in their best interest.

Marketing, medicated

Consider how this is impacting our industry: modern marketing feels like it's been over-prescribed into a stupor. Every symptom—sluggish growth, fragmented channels, lagging engagement—is treated not with insight or originality, but with another dose of technology.

What once was a discipline rooted in human understanding and long-term brand building has become, in many corners, an algorithmic arms race. Marketers are no longer interested in persuading real people—instead they're optimizing for machines: SEO crawlers, recommendation engines, social algorithms and automated bidding platforms.

The underlying belief? Every message must be measured, every campaign automated and every interaction tracked. But in trying to turn marketing into a perfectible science, we've drained it of its soul. We've taught a generation to A/B test their way to mediocrity instead of investing in bold ideas, distinctive creative and hard-earned trust.

Generative AI has only accelerated this slide. With just a few prompts, we can now churn out thousands of words or replicate brand elements across myriad channels. But at what cost? Media saturated with a flood of low-quality sameness that trains audiences to ignore everything.

[1] "Jamie Dimon has a solution to the skills shortage," Alice Tecotzky, *Business Insider*

In our rush to innovate, we've forgotten that most people don't want to just "engage with content." They want to connect with something that provides real meaning. Even performance metrics have become ends unto themselves. Clicks, impressions, open rates—easy to measure, but with little to say about whether someone cared.

There was a time when marketing introduced the world to ideas that changed culture: "Think Different," "Got Milk?," "Just Do It." Now we ship optimized banner ads no one remembers. In trying to leverage new technologies to better understand and reach people, we've made marketing less human than ever.

Tech that solves the wrong problems

This trend isn't limited to AI or marketing. Consider just a few examples from the past two decades:

- Social media algorithms: Promised connection, but delivered addiction, polarization and behavioral disorders—especially among the young.
- Smart homes and IoT devices: Marginal convenience for maximum surveillance, data risk and subscription dependency.
- Delivery drones and sidewalk robots: Gimmicks that satiate hyper-convenience culture, but do little for urban planning or local economies.
- Cryptocurrencies and NFTs: Branded as democratizing finance, they've largely enriched speculators, enabled money laundering and consumed huge amounts of energy.

Compare this to the inventions of the 19th and 20th centuries: the light bulb, telephone, penicillin, washing machines, airplanes and the personal computer. These technologies solved real problems from hunger and disease to isolation and physical labor. They gave ordinary people more autonomy, not less.

Even early internet innovation—email, open-source software, Wikipedia—was built on ideals of open access. Today's tech ecosystem feels extractive by comparison, designed more to enclose than to expand, to take more than it gives.

In our rush to innovate, we've forgotten that most people don't want to just "engage with content." They want to connect with something that provides real meaning.

The problem with "move fast and break things"

Much of today's tech harm isn't malicious—it's just reckless. Unfortunately, the "move fast and break things" mantra didn't stay in Silicon Valley. It has become a cultural ethos. Products launch before their consequences are understood. Data is collected without consent. AI platforms are being introduced into education, hiring and criminal justice systems without transparency, oversight or accountability.

We are essentially beta-testing the effects of hyperactive change and social rewiring on ourselves.

Progress should be in the service of humanity. But in a world dominated by venture capital, quarterly earnings and social media hype, the goal is no longer about building a flourishing society—it's too focused on monetization, engagement and control.

It's up to us to flip the script

The good news? It doesn't have to be this way. Technology can still serve us—but only if those of us on the front lines in our respective professions reset and reaffirm our values:

- Solve real problems, not fabricated ones
- Build slowly, ethically and with intent
- Prioritize public good over private gain
- Accept that not all friction is bad—some is essential for growth

We're at an inflection point. The last few decades have shown what happens when technological advancement outpaces wisdom. Let's be honest—we're not going to stop the Altmans and Zuckerbergs of the world. But we can demand better. We can use these new technologies smarter and we can ask harder questions—of our industry, our clients, our partners and ourselves.

That's what this issue of *Authentic Insights* is about.

The stories, strategies and perspectives that follow aren't meant to dazzle or distract. They're presented to test assumptions, surface hard truths and demand accountability. They explore how creativity, communication and technology can still work together. Not to sedate, but to energize. Not to pacify, but to empower.

If the last 20 years were defined by impetuous innovation, then let the next 20 be driven by thoughtful action. Because while we can't control the pace of technological change, we can control what we value, what we build and how we choose to show up for each other. ✧

Brian Lydon
Executive Creative Director

LYDON

Amplifying AI impact while safeguarding trust

By Marc C. Whitt, Director of Media & Strategic Relations, University of Kentucky

Let's be clear: AI won't replace the heart, empathy and critical thinking that define good public relations. What it can do is amplify our capacity to deliver timely, targeted and impactful communications.

Consider the following ways we're already seeing this in action:

- **Content ideation and drafting –** AI can help you brainstorm campaign slogans, write first drafts of blog posts or suggest social media captions.
- **Media monitoring and analysis –** AI-powered tools can scan thousands of news stories, blog posts and social media updates in seconds, spotting trends and potential issues before they hit your desk.
- **Audience insights –** AI can segment audiences more precisely and help predict what messages will resonate, allowing for more personalized outreach.
- **Repetitive task automation –** Scheduling posts, sorting contact lists, or summarizing meeting notes—AI can take those off your plate so you can focus on strategy and relationship building.

For nonprofit communicators or smaller firms, these capabilities can level the playing field, enabling small teams to compete with larger, better-resourced organizations.

The importance of ethical guardrails

Yet with all this promise comes responsibility. Just because AI *can* do something doesn't mean it *should.* PR has always been about trust—between our organizations and their stakeholders. If we abuse AI, that trust can evaporate faster than a viral rumor. Here are a few ethical considerations I outline in my recently released third book, *TAKEAWAYS: Ideas, Strategies and Encouragement for the Nonprofit Public Relations Professional:*

- **Accuracy over speed –** AI can "hallucinate"—a polite way of saying it sometimes makes things up. That's why human fact-checking is non-negotiable. We must verify every claim, statistic and attribution before it goes public.
- **Transparency with stakeholders –** If AI plays a role in developing a piece of content—especially something substantial—consider disclosing that fact. People value honesty and, in some contexts, not revealing AI involvement could be misleading.
- **Guarding against bias –** AI models learn from data that may reflect existing societal biases. Without oversight, those biases can slip into our messaging and inadvertently offend or misrepresent. We must critically assess AI-generated output to ensure it aligns with our DEI commitments and brand values.
- **Protecting privacy –** When feeding AI sensitive information—donor lists, client stories, internal memos—be sure you're not violating regulations or privacy agreements. Always know how your AI tool stores and uses data.
- **Avoiding dependency –** AI is a helpful collaborator, not your creative replacement. If you lean too heavily on it, your messaging risks sounding generic, losing the unique voice and authenticity your audience expects.

Finding the right balance

The key is to think of AI as a partner that works under your guidance. You are still the strategist, the ethical compass and the relationship-builder. AI can suggest a route, but you decide the destination.

I've had nonprofit colleagues tell me AI "feels impersonal" or is "too mechanical" for storytelling. That's a fair concern—but remember, AI is only the starting point. You bring the human touch that infuses communications with heart, nuance and credibility.

With all this promise comes responsibility. Just because AI *can* do something doesn't mean it *should.*

When I was drafting the chapter in *TAKEAWAYS* on incorporating AI into nonprofit PR operations, I kept circling back to this thought: the organizations that succeed with AI will be those that integrate it strategically, ethically, and transparently—never sacrificing trust for convenience.

In many ways, AI is simply the latest in a long line of tools that changed how we work. We adapted to email, social media and real-time analytics. We can adapt to this, too, if we stay anchored in the values that have always defined good PR: truth, fairness, respect and service to the public.

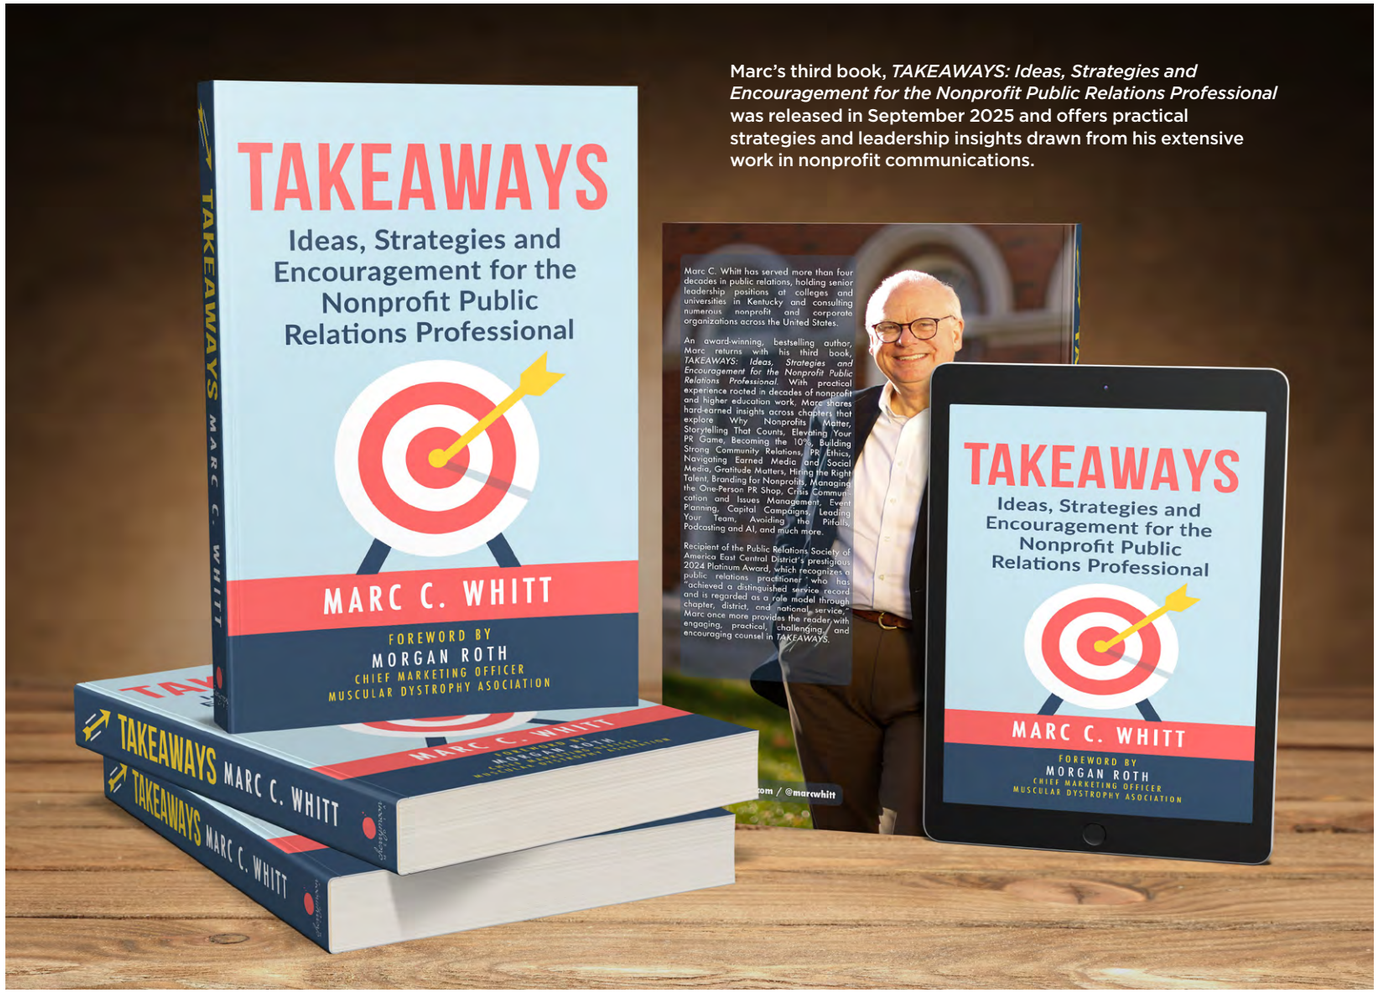

Marc's third book, *TAKEAWAYS: Ideas, Strategies and Encouragement for the Nonprofit Public Relations Professional* was released in September 2025 and offers practical strategies and leadership insights drawn from his extensive work in nonprofit communications.

Takeaways

- Use AI to enhance, not replace, human creativity and judgment. Let it handle the heavy lifting of data analysis and first drafts so you can focus on strategy and storytelling.
- Build ethical checkpoints into your workflow. Fact-check AI output, assess for bias and ensure transparency in its use.
- Protect trust at all costs. In PR—especially in the nonprofit world—credibility is your most valuable asset. No technology is worth risking it.

About the author: Marc C. Whitt is an award-winning public relations professional, bestselling author, and student-focused educator with four decades of experience. He has held senior leadership and teaching positions at public and private colleges and universities in Kentucky, while consulting for numerous corporate and nonprofit organizations across the U.S.

Rocking AI search with a “Triple S” framework

C | P

By Chris Phenner, VP of Business Development at Flash
MBA, Columbia University & U.C. Berkeley
8th Napster Employee, 15 yrs Digital Music

In the 90s, we argued about search engines the way some argued about bands. I was an early employee at RollingStone.com and then at Napster. They were music portals. Entry points for finding music content, or for music files. The internet was being reshaped by portals. Disney famously aggregated several of its brands (including ESPN) into Go.com to create The Go Network. Yahoo! was worth $125 billion in January 2000. Wall Street thought the “Search Portals” rocked.

The Search Portals weren’t like apps or websites. There was no mobile internet, so your browser’s homepage said who you were. These portals provided free email, extending their domains into the (only) messages we sent. They let us customize our homepages, making them like t-shirts worn from our favorite rock shows.

AltaVista, Ask Jeeves, Excite, Lycos and Yahoo! were The Search Portals. As with bands, nobody had data to measure why they preferred one over another. Each portal rocked for different reasons. Ask Jeeves was your butler. Excite and Yahoo! offered custom homepages. Lycos helped you build websites. Portals provided digital evidence you had graduated from AOL. Web 1.0 users also thought the Search Portals rocked.

Google arrived like Taylor Swift in her country debut, innocent and sweet. Its founders were adorable computer science PhDs from Stanford. Their PageRank system helped “algorithm” become part of the parlance of our time. From 1999 to 2010, Google climbed from zero to 90 percent market share, and it stayed there.

Until October 2024, when that all changed.

It was around then that a longtime friend—a former Apple VP I worked with at Napster—said over drinks, “I haven’t used Google in 18 months.” His declaration may have been part virtue signal, but it also points to a sea change that’s underway. Not only was he getting higher quality results via AI, but he was making a clean break from legacy search—as if he was graduating from AOL all over again.

Now, AI Search is our new scene. For two-plus years, I’ve been telling friends, “Each of us needs to re-evaluate our relationship with search for the first time in 25 years.” This is often met by blank stares (though fewer of them recently). With that in mind, I’m sharing a test I recently did and you can see if you’re hip to the new AI Search scene or learning about it from your children… or (worse) on the golf course.

AI Search requires evaluation of its corresponding platform in two parts: the first is the box into which you type your query or direction, what we now call “prompts,” and the second is the plain-language result that we now call “answers” (not the previous search results wall of blue links to which we have become accustomed).

If you type the same prompt into five or more AI platforms, you may quickly wish for the old days of Google’s blue wall. Try “Which is the best Bluetooth portable speaker?” for a brain-scrambling array of answers. Last fall, I asked five AI platforms who Kamala Harris would nominate as her running mate the day after she took over as the Democratic nominee, and several mentioned not just names that I’d never heard before but also non-existent candidates. AI hallucinations are real, and right now we are lacking criteria to help us pick the best AI platform for search.

Putting AI Search to the test

Without further ado, let’s start with a common search scenario—seeking help to determine the best use of our time and attention (our most valuable assets)—and see what we find using AI Search.

Context: I want help with an important decision about how I might spend my free time in the real world, to not only save brain cycles thinking about it but also to provide confidence of my decision in advance. With that in mind, I draft my prompt.

Prompt: “Will the new Tron movie be worth seeing in the theater?”

Google’s AI Overview (which now sits atop the legacy Google Search blue link wall) said, “It *looks like* it will be worth seeing,”—nothing wrong with that answer, other than the ten-plus seconds it took to be generated—then provided six supporting bullet points followed by a “Dive deeper into AI Mode” button. So, of course, I dove.

Google’s AI Mode hedged compared to its first cousin, saying “It *might be* worth seeing,” then provided four more supporting bullet points and three “potential downsides.” This answer ran more than 375 words and its biggest “downside” was how long it took to appear.

Google Gemini 2.5 Pro said, “All signs point to yes,” with five supporting bullet points and three additional notes to consider. I didn’t use a stopwatch, but it seemed to take much longer than its second cousins above. Comparing and understanding Google’s “AI Overview” versus “AI Mode” versus “Gemini” required three browser tabs and the juice didn’t seem to justify the squeeze.

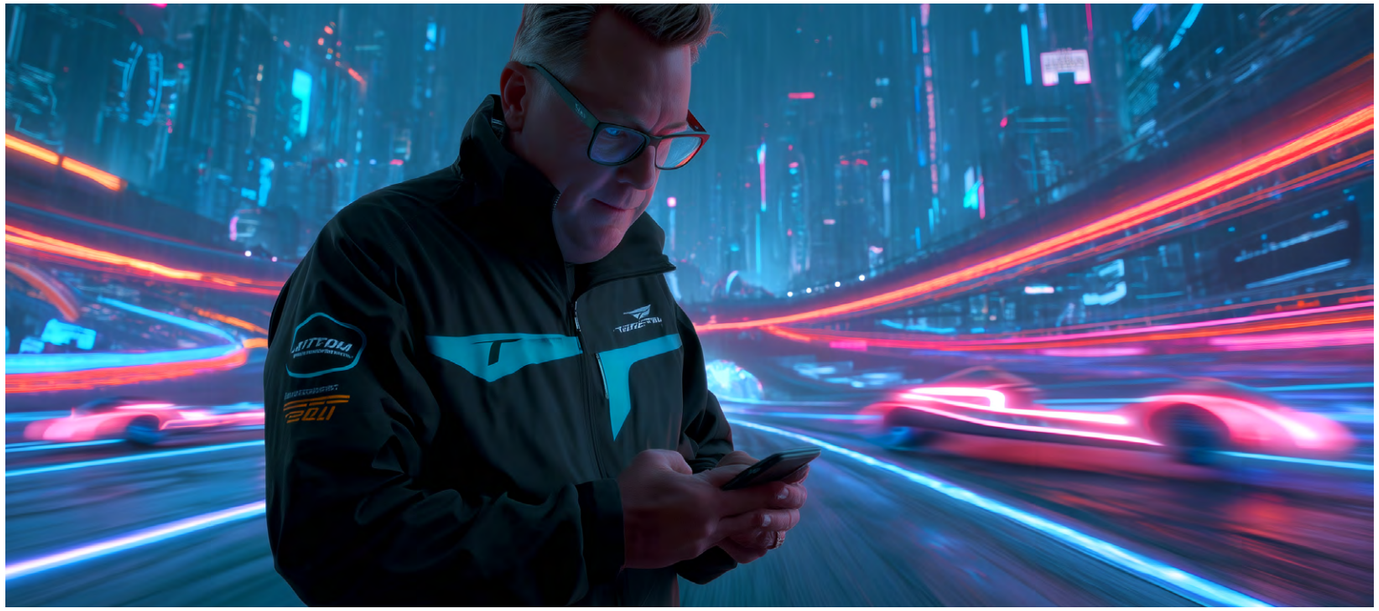

Perplexity Pro got to work immediately and delivered much faster answers than Google. Perplexity delivered paragraphs (instead of bullet points), and did not provide an "opinion" based on its findings. Ultimately, "waiting until final reviews" was its cautionary guidance. Unlike the Google AI offerings above, Perplexity provides embedded links to its sources, which has long been a unique differentiator.

ChatGPT was the second fastest to deliver its reply, which said, "The short answer is yes,"—the most definitive of all. It then supported its opinion with "Consider Before You Go" and "Final Take" sections, almost in an attempt to convince itself to attend. Attribution links to sources were provided.

Grok offered "Looks like a solid bet," for its opener (good spine!), but then spent eight lines summarizing the film—which did not help in getting to an answer. It felt like it was buying time before finally hedging with, "If you loved Tron: Legacy…" (the prior film), which ultimately avoided the original question. Three sources were cited (meh) and Grok finished with, "You might want to wait for more reviews." (So much for spine!)

Microsoft Bing may feel like a long-forgotten option, but its AI Summary (which also now sits atop its traditional search results) arrived noticeably faster than all preceding experiences. What I liked best was its statement declaring, "Generally considered worth seeing in theatres," and that its main answer page included the most links to third-party sites, containing valuable resources such as trailers and reviews.

"Triple S" framework: Speed, sources and spine

What separated AI Search that rocked from those that felt like AOL experiences of yore? I've built my rating system on three criteria:

1. **Speed:** This should speak for itself, but obviously the faster the AI platform can deliver relevant results, the more likely you are to reach for it first. Especially if you're in the lobby and movie is about to start.

2. **Sources:** I cannot imagine relying on AI answers, at least at this point in time given its known struggles with hallucinations, without valid and verifiable citation links.

3. **Spine:** In my opinion, this is what earns the most respect—when AI doesn't just provide an answer to your question, but picks a lane and backs it up.

The AI Search scene is releasing new models faster than indie bands drop new records and so far we have no clear leader—just a bunch of prompt boxes and AI platforms hoping to be considered part of the new generation of The Search Portals.

So, I encourage you to keep your current favorite AI Search platform accountable for its answers, by regularly comparing and contrasting it with others. The same simple prompt typed into a few separate AI Search boxes will help you find which one has Spine quickly—and that accountability might compel all AI platforms to continue to improve in ways that are meaningful to users. But until then, respect the process and enjoy the search. ✧

About the author: Chris Phenner has been working in internet- and tech-based roles for nearly 30 years, with the first half of his career devoted to digital music. He was the eighth employee at Napster and among the first 25 hires at both RollingStone.com and Thumbplay (later acquired by iHeartMedia). Today, he works in parking technology at Flash and likes to remind people, "Before anybody rocks out at a live show, they have to park." His other favorite cocktail-party topics include location data, mapping and partnerships.

Reports from the field

From developing websites or sharing traffic reports with a client to segmenting nurture stream audiences, there's a huge range of tasks our team handles throughout the day here at Lydon. As we continue to lean into and integrate AI across these tasks, we are collecting hundreds of amazing, frustrating, electrifying, confounding and fascinating experiences along the way.

Here are just a few impressions—the good, the bad and the ugly—from across the team. We graded the performance of the AI on various tasks, then rated ease of use on a scale of 1 (malfunction) to 10 (near perfection).

SCORECARD:
ChatGPT (Open AI)

Objective: Simulate a first-click test on a website

Summary: Wanted to see if ChatGPT (4o) could stand in for early usability testing by simulating where a first-time website visitor would click. In theory, this could save hours of setup and give quick directional insights into navigation clarity before investing in formal research.

Prompting: I began by asking ChatGPT how to simulate a first-click usability test. It recommended uploading a homepage screenshot and offered a sample prompt. So far, so good.

I uploaded the image and entered: "Here is a screenshot of [company's] homepage. You're a [target user] looking for help with [service]. Imagine you are a first-time user of this website. You want to [goal]. Where would you click first and why?"

Rather than simulate a user's behavior or provide rationale, ChatGPT responded with instructions for how I should navigate the site—completely missing the point of the exercise. I tried again, this time clarifying: "Take the perspective of a UX researcher. Where would a [target user] most likely click first—and why?"

Result: At first, the feedback looked more promising. ChatGPT hypothesized a likely first click (the "Explore our Solutions" section) and provided a rationale tied to semantic alignment between the user's goal and that label. It even offered a secondary click option (the "What We Do" section), explaining how it might perform better in an eye-tracking or heatmap test.

The problem? Neither section actually existed on the site. ChatGPT had fabricated entire clickable areas—labels, structure and navigation paths—presenting them with confidence as if they were real. Instead of surfacing potential usability issues, I ended up chasing phantom insights. What should have been a quick test devolved into a half-hour of verifying hallucinations before I finally moved on.

The exercise highlights a deeper limitation: generative AI can mimic the language of research, but without grounding in actual artifacts (like a functioning site map), it risks producing output that looks authoritative yet collapses under scrutiny. Treating it as a heuristic reviewer might someday save time—but today, it's more illusion than insight.

Grade: F

Ease of use: 0

– C. Moser

SCORECARD:
LetsEnhance

Objective: Improve the quality of photography taken with an iPhone at a live event for future promotional use

Summary: The test was whether LetsEnhance—a subscription-based platform that promises to upscale photos by boosting detail, resolution, and overall "professionalism"—could take noisy, low-light iPhone event photos and make them look like polished, brand-ready shots without requiring expensive cameras or reshoots.

Prompting: The LetsEnhance workflow is simple: upload, adjust settings and download a higher-quality version.

Result: Backgrounds and textures improved dramatically—stage lighting, signage and crowd details looked sharper and more usable. But people were another story: faces crossed the uncanny valley instantly, taking on that plasticky, overprocessed "AI look." The only salvageable workflow was hybrid—keeping the original faces and compositing them back into the enhanced frame. That trick rescued about 90 percent of the photo, but it was extra effort. The bottom line: good for venue shots, risky for human subjects.

Grade: C+

Ease of use: 6

– R. Hock

SCORECARD:
ChatGPT (Open AI)

Objective: Learn to build a Notion workspace to serve as a single source of truth for projects, tasks, and related ideas.

Summary: Could ChatGPT double as a Notion consultant—helping design not just a workspace, but a scalable system for managing work?

Prompt: I asked ChatGPT to outline a workspace structure and provide clear setup instructions.

Result: The AI delivered a surprisingly clear, beginner-friendly walkthrough. Its step-by-step guidance let me configure a working Notion hub in a fraction of the time it would've taken on my own. Even as a first-timer, I could follow along easily. But here's the catch: the output stopped at a "starter template" level. It didn't anticipate advanced cases like database roll-ups, permissions or cross-team scalability—the features that separate hobbyist setups from enterprise-ready systems. My takeaway: ChatGPT is excellent at accelerating the grunt work of setup, but it doesn't yet act like a true architect. For more sophisticated use, I'd want it to not just explain how to set up blocks, but to ask probing questions about workflow, governance and failure points—like a real productivity consultant would.

Grade: B

Ease of use: 7.5

– K. Bogott

SCORECARD:

Grok (X)

Objective: Generate multiple headlines for articles for a financial services client

Summary: The experiment tested whether Grok (3) could move beyond functional rephrasing to deliver headlines that actually sell ideas—where word choice carries weight with a discerning professional audience.

Prompt: After uploading four human-drafted thought leadership articles with audience context and personas, I asked Grok to generate alternate headlines.

Result: Grok produced headlines that were technically correct but creatively hollow. They leaned on safe clichés and generic phrasing, ignoring the subtle cues in the articles that could have been leveraged for authority or intrigue. In other words, it could shuffle the words but not sharpen the hook. To get closer, I had to repeatedly reframe prompts with explicit instructions on tone, competitive positioning, and rhetorical angle. That narrowed the gap—but it still didn't deliver the kind of persuasive specificity financial services audiences expect. Ultimately, I had to rewrite each headline myself, cannibalizing only fragments from Grok's drafts. I learned that AI isn't yet a headline generator—it's a brainstorming mirror. Useful for volume and variants, but not for the editorial instinct that makes one line worth clicking in a sea of sameness.

Grade: C-

Ease of use: 4

– M. Stefanowitz

SCORECARD:

Sora (Open AI)

Objective: Generate a short promo video for an IT Services client

Summary: Tested whether Sora could create a 20-second promo—something simple but polished enough to showcase a client's B2B enterprise technology services on LinkedIn without requiring a shoot, stock video or heavy post-production.

Prompting: "A professional promo for an IT services firm. Office environment, shots of diverse teams collaborating, screens with data visualizations, a confident executive presenting to clients. Modern, credible, business-focused tone."

Result: At first glance, the clips looked promising—but closer inspection revealed major flaws in realism. I refined my prompts, but each iteration introduced new distortions. In B2B marketing professionalism and precision are non-negotiable. The only way forward would have been heavy editing or compositing, but even then the footage wouldn't have met our standards required for client-facing deliverables. It's clear that gen AI can spark ideas—but instead of saving production time, it creates unpredictability making deadlines impossible to manage. For now, it's better as a tool for pitching concepts than producing actual client-ready assets.

Grade: D

Ease of use: 2

– K. Meyers

SCORECARD:

ChatGPT (Open AI) vs. Claude (Anthropic)

Objective: Compare ChatGPT Custom GPTs vs. Claude Projects for strategic growth planning

Summary: ChatGPT's "Custom GPTs" feature and Claude's "Projects" feature promise a solution to one of Gen AI's biggest weaknesses: memory. Both can provide some degree of contextual knowledge retention via the use of a "Memory Pack"—simply, a set of reference documents that can be uploaded and stored in a Custom GPT and Claude Project. By uploading Lydon's business plan, branding and identity system, marketing content and strategic internal workflows, I wanted to see if either tool could not only accurately retain and reference specific context, but also generate fresh, strategic thinking about Lydon's future growth. I named the assistant "Lydon Agency Growth Partner" (LAGP) in both platforms.

Prompting: Once the Memory Pack was in place, I queried my LAGP in both ChatGPT and Claude with the same prompt: "Propose a new growth strategy for Lydon based on our capabilities." ChatGPT proposed evolving from project-based work into a platform-first model, emphasizing scalable solutions, repeatable categories, and deeper client relationships. Claude responded with a similar structure, highlighting our Marketing Platforms as the core differentiator and breaking growth into pillars, metrics and implementation.

The problem was that both felt like echoes of the Memory Pack. Useful, yes—but more like a strategist paraphrasing a client brief than generating true breakthrough thinking. So, I followed up with: "You're just repeating back my Memory Pack. Give me something novel yet realistic."

Result: Here I saw some divergence in the tools' thinking.

ChatGPT pitched "Lydon Intelligence Networks"—positioning the agency as a network orchestrator that connects clients, industry communities, and insight-sharing ecosystems. It was ambitious and visionary, if a bit abstract: imagine Lydon as part agency, part research bureau, part publishing platform.

Claude proposed the "Marketing Intelligence Engine"—a consultancy model that monetizes proprietary market intelligence streams. More grounded, more revenue-focused and closer to something that could actually be packaged and monetized.

In the end, both platforms demonstrated the power of memory-enabled AI to structure strategy quickly, but also their tendency to recycle what's already there. When pushed, ChatGPT seemed to lean visionary but abstract while Claude leaned more practical and monetizable. The real potential lies in how these tools force sharper thinking—not because they hand you the answer, but because they challenge you to interrogate your own strategy through a new lens. I'll definitely be exploring their use further—and, who knows, you might see a new service offering, value proposition, or brand repositioning from Lydon sometime in the near future.

Grade: B for Claude, B- for ChatGPT

Ease of use: 7 (for both)

– B. Lydon

Prime prompting: Ask (the right way) and you shall receive

By Sean Griffin, Strategy & Content Director

Most people think better AI prompting means longer prompts. Or more precise ones. Or a massive library of copy-paste templates. In reality, the most effective prompting strategies are often the simplest—not because they're clever, but because they're clear.

As a professional writer with a degree in a second language, I've spent years thinking about how phrasing shapes outcomes—how small changes in tone, framing, or order can shift the way something lands. And in using AI across different tools and tasks, I've noticed that some of the same principles still hold: clarity over cleverness, pacing over pressure and intention over length.

Large language models don't read between the lines. They don't infer intent and they don't pause to ask clarifying questions. So what you say—and how you say it—matters. A lot.

Below are five subtle but powerful shifts that consistently lead to better, faster results. These aren't just tips for getting cleaner copy. They're ways to think more clearly in the prompt itself—no matter what kind of output you're working toward.

1) Be specific, not soft

A lot of prompts start with hedging: "Make this a bit more friendly," "See if that can be little tighter," or "Don't be too casual." Many of us are conditioned to make requests of others collaborative in tone. That's nice for a human, but to a LLM it's a vague instruction—and the output you get back will reflect that. To improve your results, be precise. That requires anchoring your request in a real-world tone, voice or scenario, and giving it a clear persona or context to emulate.

Instead of: "Try to make this a bit more persuasive."

Try: "Write this like a product manager explaining the business value of a feature to a skeptical executive team."

Instead of: "Maybe just make this a little warmer?"

Try: "Write this like a welcome email from a helpful product manager—friendly and confident, not overly casual."

The more you define the tone, the less the model has to guess—and the more the output reflects your intent.

2) Use analogies that carry shape, not just style

When prompts fall flat, it's often because we only define tone—not form. We'll say "make it sound more engaging" or "add a little creativity," but leave the model to invent its own structure. Analogies can help fix that. They bring both shape and style. When you say, "Write this like a museum docent introducing an exhibit," you're offering tone, pacing, point of view and narrative intent—all in one.

Instead of: "See if you can make this more dynamic."

Try: "Frame this like a launch keynote—start bold, build intrigue and land it with a clear takeaway."

Instead of: "We want it to sound intelligent but fun."

Try: "Write this like an NPR podcast intro—thoughtful, curious and a little personal."

Analogies act as shortcuts to a shared format. They help the model channel something you both will understand—even without defining it line by line or bullet by bullet.

3) Slow the prompt, speed the progress

Most of us try to be efficient in our prompts: "Summarize the research and turn it into five blog post ideas with SEO headlines and tone notes." But cramming that much into one ask forces the model to deliver a little of everything—and not enough of anything.

It's better to break the task into steps. Ask for insights first, use them to shape directions, then refine tone or format. Each step adds structure and clarity—without creating extra work."

The following example shows how to break up a request into steps:

Step 1: “Summarize the five most useful takeaways from this research.”

Step 2: “Based on those, generate five blog post angles.”

Step 3: “Now give each a headline and a tone-of-voice label.”

The same approach works for strategy or concept reviews.

Instead of: “Review this brand strategy doc and tell me what’s missing or unclear, and then summarize the core value prop,”

Try: “First, list any areas in this strategy doc that seem inconsistent or underdeveloped,” then, “Now summarize the core value proposition in one sentence,” and finally, “Next, highlight anything that might confuse an outside reader.”

This kind of prompting may be slower, but the important thing is that it’s structured and deliberate. You’re guiding the model like you’d guide a collaborator. Each step gives you a moment to assess, steer or deepen the work. So you may be adding steps, but you’re also reducing restarts.

4) Prime before you prompt

Prompting usually starts with the ask. But when the task is nuanced—with layers of creative development, stakeholder framing, logic problems and more—it helps to start with setup, and we call that “priming.”

Priming is simple: give the model context before you make a request. Upload a brief, define the audience and frame the tone. Sometimes it’s as direct as saying, “We’re writing for frontline healthcare managers. They’re skeptical of buzzwords, short on time and want practical tools over inspirational language. I’ll share the copy next.”

Other times, you can even ask the model to help you make stronger, clearer requests: “What’s the best way to prompt you to critique the structure of a positioning statement?”

Priming gives the model more to go on. It doesn’t just improve alignment—it raises the ceiling for what the output can become.

5) Ask what the model thinks first

Most prompting is reactive: generate this, fix that or rewrite it better. But sometimes the most useful prompt isn’t a request—it’s a question. Before you ask for a deliverable, try asking the model to interpret or critique what it sees. It slows the process just enough to check alignment before the words hit the page.

Consider the following examples:

“Based on this content, what’s the central message you think I’m trying to communicate?”

“What’s missing or unclear in this argument?”

“What would a skeptical reader ask after reading this draft?”

This is different from priming. You’re not feeding the model information—you’re asking it to think. That shift from instruction to reflection can help surface gaps, sharpen your brief or challenge your assumptions before the real work begins.

Soft prompting	Specific prompting
“Can you make this a bit more engaging?”	“Frame this like a TED Talk opener—clear and idea-driven.”
“Try something more professional.”	“Use the tone of a B2B case study or VC pitch for senior execs: confident and benefit-forward.”
“Sounds boring, can you make it pop?”	“Rewrite this in a tone similar to a high-energy product launch email.”
“I need this to be more convincing.”	“Rewrite this in the voice of a nonprofit grant writer.”
“Can this be just a bit more fun?”	“Use the tone of a friendly narrator in a kids’ science video.”
“Make it sound smart.”	“Write this like a New Yorker explainer: Layered, thoughtful and precise.”
“Need this to have more of a salesy feel.”	“Make it sound like a SaaS landing page—with headline, benefit and CTA.”
“Let’s make it sound more modern.”	“Use a tone similar to Wired: crisp and future-facing.”
“Keep it conversational.”	“Write it like a peer-to-peer Slack message—casual but smart.”
“Make sure it’s not too stuffy.”	“Use a tone like Fast Company—professional and accessible.”

Remember—prompting is a skill, not a shortcut

These five strategies aren’t hacks or templates. They’re habits I’ve developed by applying what I know about phrasing, structure and intent—skills sharpened over years of writing for teams, clients, and ideas that needed clarity to land well.

Prompting well doesn’t require you to be a writer. But it does require you to think like one: to anticipate misinterpretation, to frame your ask with intention and to guide the output with precision, not just volume.

Good prompting won’t eliminate iteration. But it will make each round more useful and less frustrating. Think twice, prompt once—and get what you actually need. ✧

Vibe coding is no silver bullet but can still provide useful ammo

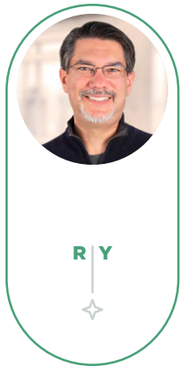

By Rick Yager, Executive Creative & Technology Director

If you had told me a year ago that I'd be spending my evenings writing and revising actual working code with an AI assistant, I would've laughed—then asked if you could help fix my JavaScript.

For most of my career, I've straddled the creative and technology sides of agency life. My background is in design, not development. Sure, I've dabbled: I survived a True BASIC class in college thanks to a frat brother and my computer science roommate. I built my first website in 1996 after a semester of evening HTML/CSS classes at a community college. And like many designers-turned-marketers, I've made a decent career slicing PSDs into templates and customizing WordPress themes. But full-stack AI SaaS platforms? That always felt like someone else's job.

That changed this year.

The shift didn't come from a grand vision—it came from friction, and a deeply ingrained D.I.Y. mentality. I had a clear idea for a solution: I could visualize the components, sketch the UX, and articulate the logic. But I couldn't build it—at least not the conventional way. The idea of investing significant time and money working with developers on a rough concept pushed me to try something different: I asked AI to help me build it.

The beginning was clunky. I didn't follow a course or take a tutorial. I opened ChatGPT and typed: "I want to build a modular tool that collects input, calls APIs, runs LLM processes, and returns formatted output to the user. What's the best architecture for that?"

To my surprise, it answered—clearly. I asked follow-ups. It adapted. I challenged its assumptions. It revised. Before long, I wasn't just using AI to explain how things should work—I was using it to configure tools and write the components. When I didn't understand something, I told it so. And it slowed down to teach me. It became a silent tutor, a pair programmer and a translator, all in one.

I didn't know it had a name when I started, but this process is now often called "vibe coding"—a new trend where non-developers use AI to build software through intuition, trial and error, and rapid iteration.

On the frontier of AI, sometimes even imperfect shots are worth taking if you can score a hit.

Every misfire is a lesson

One of the first lessons I learned? AI doesn't always give you the same answer twice. It's a pattern-matching engine, not a source of absolute truth, so its memory recall also leaves a lot to be desired. I'd ask it to solve a problem, get one approach one day and a completely different method the next—even with nearly identical prompts.

That inconsistency can be frustrating but it's the nature of LLMs at this current point in time. Over time, I've learned to enhance its memory and anticipate where things might go sideways—and how to prompt around the pitfalls. This isn't like following a recipe. It's more of a collaboration—with a virtual partner that can often be just as fallible as a human.

I found that the most critical part of this process is the feedback loop—ask the model to generate something, spot where it falls short, revise your input and try again. These loops not only teach you how to better shape your prompts, they gradually teach the AI to better align with your goals. They're often insightful, sometimes maddening, but they always move the work forward.

Vibe coding has indeed been the target of well-deserved skepticism and criticism as of late. Recent news articles and reports have shot holes in claims of increased developer productivity or profitability within software and technology companies, including one from Bain & Company[1]: "Generative AI arrived on the scene with sky-high expectations, and many companies rushed into pilot projects," the report reads. "Yet the results haven't lived up to the hype." It would seem that expert, experienced developers aren't going away anytime soon.

[1] "From Pilots to Payoff: Generative AI in Software Development," Purna Doddapaneni, Bill Radzevych, Steven Breeden, Bharat Bansal, Tanvee Rao, *Bain & Company Technology Report 2025*

Progress isn't always about hitting the bullseye

Despite my adventures into vibe coding and everything I've learned, I'm still far from a developer by traditional standards. But, I've built multiple working AI-powered modules. I've learned my way around GitHub. I've deployed frontends to Vercel. I've chained model calls, debugged formatting errors, built logic branches and implemented authentication, all with AI guiding the process.

More importantly, vibe coding has challenged me to think differently about how I utilize AI. It's not just a productivity tool. It's not even just an assistant. It's a collaborator, one that forces you to clarify your thinking and challenges you when logic breaks down. It's helped me become more precise, more patient, and oddly, a little more adventurous.

And while I hope to create a fully functional piece of software that we can use internally—and potentially bring to market—I'd still count this effort as a success even if it becomes nothing more than a well-formed proof-of-concept that we polish and scale with real developers. The important thing is, I have been able to design and build it with clarity and conviction.

For anyone curious but unsure where to begin—start with what you already know. Let AI do the heavy lifting where you're weak and don't be afraid to lean on it. It can be surprisingly helpful, occasionally frustrating, and—much like a human partner—it works best when you meet it halfway. ✧

AI can make moments—creatives make them matter

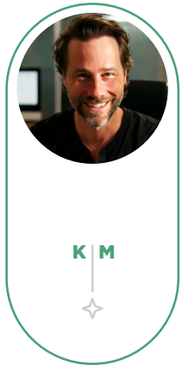

By Kris Meyers, Senior Designer, Motion & Interactive

Need a product demo? AI can show your phone gleaming on a pedestal...but in the next frame, it's upside down, on fire and inexplicably in the ocean. That's the magic—and the migraine—of generative AI video in its current state.

The pace of progress here is dizzying. Blink, and a new "best-in-class" model drops, making yesterday's jaw-dropping output look like a low-res relic. Keeping up with what's possible in generative AI applications, including Veo3, Kling, Runway and Sora is practically a full-time pursuit. And yet, no matter how real the lighting, how smooth the dolly shot or how clever the compositing, these tools are great at producing moments, but not cohesive, client-ready productions.

What AI gets right

The appeal is obvious. With the right text prompt, you can generate cinematic-quality clips in minutes, no camera crew or green screen required. Want a sweeping drone shot over a futuristic city at sunset? Veo3 will nail it with lighting and motion so convincing you'll check for a permit. Need an artfully stylized, surreal transition? Runway makes it feel like you have a film school's worth of effects artists on speed dial.

There's also the pure creative freedom, from impossible camera moves to fantasy landscapes and physics-defying action sequences. For smaller agencies or solo creators, this delivers high-end production quality without the high-end spend.

Democratizing and the designer's dilemma

I'll admit, it's a little surreal watching tools hand cinematic capabilities to anyone with Wi-Fi and a few spare minutes. On the surface, it can make seasoned designers feel like our years of technical skill are suddenly less valuable.

But here's the truth: AI video isn't replacing us. Not now, maybe not ever. What it is doing is expanding the sandbox. A motion designer can explore cinematography. A copywriter can dabble in art direction. A graphic designer can storyboard a concept and see it rendered as a living, moving sequence faster than you can boot up After Effects.

Far from making us redundant, these tools can make us faster, more experimental and more cross-disciplinary. They're power tools. But like all power tools, give them to a pro and you get precision. Give them to an amateur and you get a very creative accident.

Where AI falls apart

Ask AI video for a single shot, and you'll often get something impressive at first glance. Ask for a sequence, and the cracks show fast. Characters change outfits between frames. Props disappear. A coffee mug turns into a wine glass mid-scene. Physics occasionally stages a coup—liquids pour upward, runners sprint like marionettes in zero gravity, shadows move in ways that would terrify any lighting designer.

Even with the strongest models, getting consistency is a grind. You end up generating multiple variations, cherry-picking usable seconds and stitching them together in post just to maintain basic continuity. It's like working with an intern who delivers flashes of brilliance...but has no short-term memory.

How the sausage gets made

Concept & planning

Define the look, mood, and pacing upfront, then match tools to strengths: Veo3 and Sora for photorealism and motion; Runway and Higgsfield for stylization; Heygen for avatars; ElevenLabs for voice; Suna for music. Clear direction early ensures each tool pulls its weight.

Prompting & testing

Refine prompts for style, shots, and movement—test, tweak, repeat. Custom GPTs can help lock in vibe, lens, and aesthetic. Iteration is where quality builds, but it's also the most time- and credit-intensive stage, with each render costing both budget and patience.

Editing & assembly

Export the best clips, stitch them together, and fix continuity. Smooth transitions, color grade, and add sound in Premiere; polish with motion graphics in After Effects. This final pass ensures pacing, branding, and intent come through clearly.

Hybrid workflows, human stories

Generative video excels at producing striking, isolated moments—but humans still craft the stories. Without a creative lead, pacing falters, continuity breaks and emotional beats fail to land. The current sweet spot lies in hybrid workflows: let AI accelerate ideation, generate raw footage, and suggest visual directions you might never have considered, then let human direction and editing shape it into something coherent, compelling and client-ready.

If the past 18 months are any indication, quality will improve quickly. Models will manage continuity more reliably, physics will stabilize and iteration cycles will shrink. Yet narrative cohesion still lags behind the visual spectacle—and that's okay. The goal isn't to hand over the designer's chair, but to expand the creative toolkit. The best results will come from teams who view these tools not as replacements, but as accelerators, collaborators and idea machines.

Until AI can figure out how to get its moments to connect in a way that feels genuine and hits on an emotional level deeper than surface sheen, the story is still ours to tell. ✧

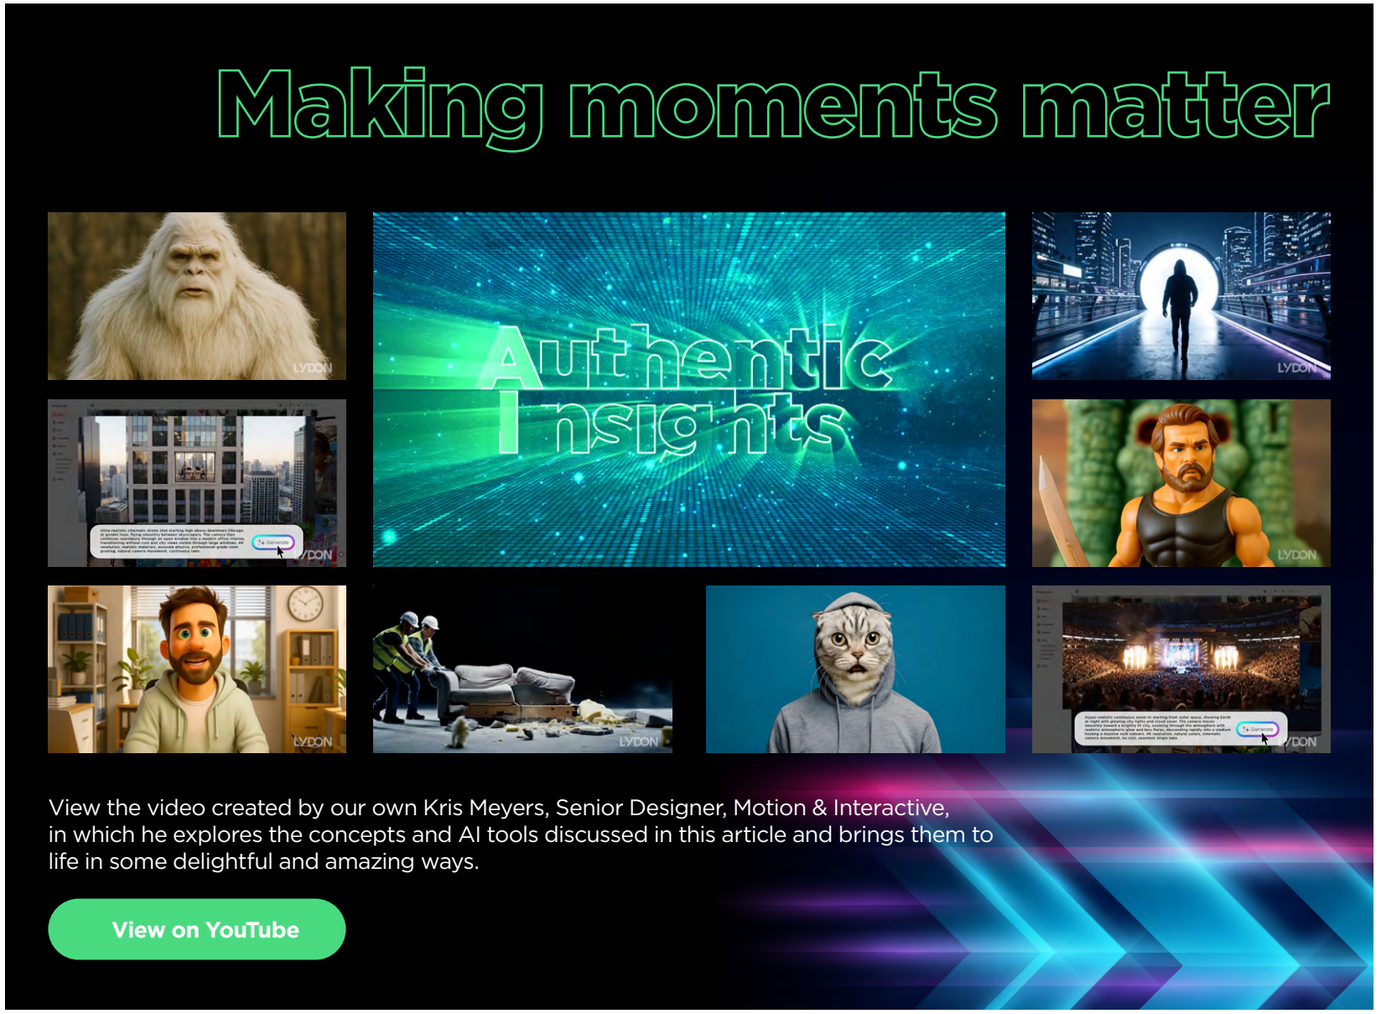

Clearing the clutter to make way for great work

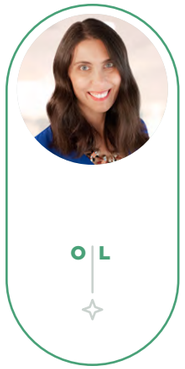

By Olivia Lydon, Client Engagement Director

As a business development leader, I often have sudden or unusual tasks for the team. It's part of my role to keep projects moving—often across departments—while client work continues on all fronts.

In the fast-paced, high-stakes world of earning new business, there's real value in how well a team operates: how easily they collaborate, how quickly they can adjust and how much time they're able to spend on work that actually moves objectives forward. Lately, I've been exploring how AI can support those dynamics in everyday ways. Not just as a creative tool, but as something that helps behind the scenes—making it easier to plan, coordinate and execute.

One platform that's become especially useful for our team is Claude. It handles messy inputs well, from research decks and campaign plans to folders full of PDFs. It doesn't need constant re-prompting, and it keeps pace with the way our teams work: fast, collaborative and often changing in real time. Here are a few examples of how we've been using it.

Cut through the chaos during campaign kickoffs

Campaign work often begins with a tangle of communications and documents—brand guidelines, creative briefs, memos, emails, reference materials and more. Claude can process all of it at once and provide a high-level summary in minutes. That's time saved on sorting and more time spent on strategic thinking.

How to use it: Upload relevant files and prompt Claude to identify key takeaways, gaps or performance insights. This small shift will quiet the noise so real thinking can begin.

Keep strategy flowing, not fragmented

Strategic thinking often gets paused for necessary, dedicated research. Claude helps keep research integrated into planning. Whether you're looking at competitor campaigns or pulling market trends, it can surface insights while you continue to build your plan.

How to use it: If browsing is enabled, prompt Claude to "find recent examples in [industry]" or "summarize current trends in [topic]."

Track changes without losing the thread

Plans and briefs change, sometimes rapidly. Claude helps keep those documents updated without losing alignment. Teams can collaborate while maintaining version control and avoiding long email chains and disjointed IM threads.

How to use it: Upload a working document and ask Claude to maintain structure while tracking edits or integrating new information.

Make the data work harder so you can work smarter

Turning raw data into something you can actually use—like a deck or a summary—can slow down a good idea. Claude can generate clean outputs from messy inputs, helping you move faster.

How to use it: Provide performance data or paste in a CSV. Then prompt Claude to generate charts, slides or a concise summary based on the uploads. When data isn't a barrier, the team can stay focused on what they're building and spend less time on the reporting.

Bring ideas off the page sooner

When a mockup isn't enough to tell the story, Claude can help you create something interactive for early feedback. It's a way to get teams aligned quickly and reduce misfires in the next round.

How to use it: Upload a layout or wireframe and ask Claude to suggest a UI flow or generate a basic coded prototype.

When to use tools like Claude

In our experience, AI is most valuable when it's used with intent. It's not always a shortcut, but a support system—especially helpful when you're sorting complexity or trying to build momentum.

It won't replace your creative instincts or decision-making. But it can free up more time for both. Whether you're reviewing a campaign, shaping a brief, or getting early feedback on a concept, tools like Claude can help you spend less time untangling the process and more time moving forward. The advantage isn't just the technology. It's knowing where it fits and how to make it work for your team.

That's what it means to clear the clutter to make way for great work.

A working AI glossary

Terms marketers are hearing (and hesitant to ask about)

AI slop: A slang term for bland, error-filled or overly confident Gen AI output—especially when it's used without editing or human oversight.

Relevancy for you: You've probably already seen it in presentations, emails or social media, and it's a key reason why human review remains critical.

Autonomous agents: Also known as "Agentic AI" these systems will determine and act on a series of steps that can manipulate other systems without constant human prompting or intervention.

Relevancy for you: Still an emerging technology, but it's being heavily discussed in marketing automation and operations for its potential to streamline complex workflows.

Chain-of-thought prompting: A method of prompting where you instruct a model to reason through a problem step-by-step, rather than jumping directly to a final answer.

Relevancy for you: This technique often produces more accurate and useful output, especially for complex or multi-step tasks.

Context window: The amount of information that an AI model can consider at one time, including the prompt itself and any previous conversation history.

Relevancy for you: This determines whether a lengthy prompt, a large document or an entire conversation will be fully considered by the model.

Embedding: A way of converting concepts (like words, images or entire documents) into numerical representations that an AI model can use to compare and group them.

Relevancy for you: This is a core technology behind search, personalization and similarity tools, though it's typically abstracted away from the user.

Fine-tuning: Training an AI model further on a specific dataset to improve its performance on a niche task or to align its responses with a particular brand voice.

Relevancy for you: It's a common approach for building brand- or product-specific AI assistants and content tools.

Guardrails: Mechanisms that set boundaries on what an AI tool can say or do. They are designed to reduce risk, prevent harmful outputs or ensure compliance and on-brand behavior.

Relevancy for you: This is a key feature highlighted by vendors when they claim to have "safe" or "brand-aligned" AI tools.

Hallucination: When an AI model confidently generates false, made-up or nonsensical information.

Relevancy for you: If left unchecked, this can lead to major trust issues in your content, research and communications.

Inference: The process of running a trained AI model to generate an output based on a new input.

Relevancy for you: As the computational process behind every AI response, it is sometimes used to explain slow load times or usage limits.

Latent space: The internal "map" or abstract representation that an AI model builds to understand the relationships between different data points and concepts.

Relevancy for you: While not essential to understand in detail, it's a term that often appears in vendor presentations and AI thought pieces to describe how models "think."

Multi-modal: An AI system that can interpret and generate across more than one type of data, such as text, images, video, or audio.

Relevancy for you: This capability enables tools that can generate a presentation from a text prompt or create a storyboard from a series of images and descriptions.

Prompt engineering: The burgeoning craft of writing effective prompts to guide an AI model's behavior and output toward a desired result.

Relevancy for you: This skill is often the single biggest difference between getting a generic answer and a genuinely useful or insightful result from an AI model.

Prompt fatigue: A human limitation where a user becomes mentally tired from constantly guiding AI tools with trial-and-error prompts.

Relevancy for you: This is a real friction point that can affect both the adoption of AI tools and user creativity.

RAG (retrieval-augmented generation): A technique that combines a language model with a search over a specific, external database or document set to give more grounded and factual responses.

Relevancy for you: This is the technology powering more custom, "enterprise-aware" AI tools that is capable of accessing and using your proprietary information.

Synthetic data: Artificially generated data used to train AI models or simulate audience behavior.

Relevancy for you: It's being used in some marketing research, segmentation and testing tools to create realistic scenarios without using actual customer data.

Tokens: The building blocks that AI models use to process language. They are typically chunks of words, characters or subwords.

Relevancy for you: This is the primary unit that affects pricing, length limits and prompt formatting in many AI tools.

Vector database: A database designed to efficiently store and search for embeddings. It is often paired with RAG systems.

Relevancy for you: This technology powers "smart" search functions in AI-enhanced knowledge bases and content management systems.

Zero-shot and few-shot prompting: Getting a model to complete a task with no examples (zero-shot) or a minimal number of examples (few-shot).

Relevancy for you: This impacts how much setup and example-based learning are needed to get high-quality results from a model.

Why AI will probably never get creativity

By Sean Griffin, Strategy & Content Director

Let's get something out of the way up front. AI is fast. It's efficient. It can draft a decent headline, rewrite a paragraph or pull together a tidy SEO listicle in under 30 seconds. But, when it comes to *generating* something original—something timely, relevant and emotionally attuned—generative AI still comes up short.

That gap matters. And it's exactly where creative agencies are proving their ongoing value.

Recently, I read a reflection from Scott Adams, the creator of the *Dilbert* comic strip, who asked ChatGPT to help him come up with ideas for the day's strip. Not even the punchline, just the setup. The AI responded with a list of ideas that read pretty much how you'd expect: the boss doesn't give you a raise, a coworker takes credit for your work, someone emails a spreadsheet at 4:59 p.m. on a Friday. None of it was wrong, technically. But it also wasn't original, interesting or particularly amusing.

This is the central limitation of generative AI in creative settings: it can't feel what people are feeling. It doesn't have a pulse on culture, mood or tension in the moment. It looks backward. It recombines what's been said and done. What it produces is technically accurate—but often emotionally empty. And that's a problem if you're trying to reach real people.

Creative professionals—especially those inside multidisciplinary agencies—operate differently. When we sit down to solve a problem, we're not just referencing the past. We're scanning what's happening right now: the news, the market, the cultural temperature. We're paying attention to nuance. We're asking, what are people feeling that they don't normally feel? And how can we respond to that?

It's this sense of emotional context and narrative timing that makes creative work resonate—and it's where agencies deliver value that no off-the-shelf AI can.

That doesn't mean we ignore the tools. At Lydon, like many teams, we're using AI to accelerate production tasks, support ideation, and draft early versions of copy. We're testing it constantly. But we're also clear-eyed about its (very real) limitations:

- It doesn't know your audience on a human level the way we do.
- It can't offer a fresh point of view when the market is chaotic and changing rapidly.
- It can't lead with empathy or tone that feels just right.
- And it certainly doesn't make creative leaps based on something no one's said out loud yet—but what your intuition tells you that everyone is thinking.

That's what strategy and creative teams are built for.

We don't claim exclusivity over insight. But we do claim process: the ability to sit with ambiguity, test a dozen angles and pursue the one that feels the most human. That's not "inefficiency" as some might have you believe—that's actually where the magic lives.

There's a temptation right now, especially in fast-moving organizations, to see creativity as something that can be automated. But creativity is not just a series of inputs and outputs. It's a conversation. A question asked the right way. A hunch. It happens in the hallway, on the commute, in the shower. It's the flash of brilliance that strikes in the space between the brief and the brainstorm.

Creativity is not just a series of inputs and outputs. It's the flash of brilliance that strikes in the space between the brief and the brainstorm.

That's what good agencies offer: the space, structure and focus to do this kind of work well. And, increasingly, that work isn't just "nice to have," it's what will make marketing feel necessary again.

We're not here to beat AI. We're here to leverage it—but also to remain rooted in the one thing it still can't replicate: how it feels to be human and possess the intuition to know what might move someone, right now, to care about your mission and your message.

Content

All copy in this issue was written entirely by the *human* authors as indicated on each article, with the following considerations:

The pernicious promise of progress (p.3) – Had various sections "punched up" by Chat GPT and some stats researched. Author was responsible for final draft and validating all stats.

Prime prompting: Ask (the right way) and you shall receive (p.12) – ChatGPT generated examples for each prompting concept written up by the author, who then filtered and selected those he felt were most helpful for this audience. Final draft was the work of the author.

Vibe coding is no silver bullet but can still provide useful ammo (p.14) – After finishing coding work with the help of ChatGPT, asked it to draft an outline for an article about the experience in the same chat, followed by suggested refinements. Final draft was the work of the author.

AI can make moments—creatives make them matter (p.17) – ChatGPT was used to spark ideas and structure a draft, but the final writing, editing and creative decisions were made by the author.

Clearing the clutter to make way for great work (p.19) – Claude was asked for ideas on how it would streamlined processes across Lydon's dynamic marketing team. The author then used ChatGPT to suggest real-world applications. Final draft was the work of the author.

Why AI will probably never get creativity (p.19) – Started with the real-world anecdote mentioned in the article, discussed the concept with ChatGPT (with some agreements and several arguments), then distilled and heavily refined the best ideas into a single, focused narrative. Final draft was the work of the author.

Cover image

Origin: Artificial

Technique: Started with an authentic image search of a cowboy rustling cattle (chasing an idea of the "AI frontier" called out in our tagline), then, based on one of the shots I saw, I got an idea I knew was better suited for Gen AI. Using Midjourney, and the following prompt, I got it within a few attempts: "A cinematic shot of an old cowboy riding his horse through the tall grass, but with a quirky and clunky little steampunk robot strolling alongside him. The setting is moody and atmospheric, with low-hanging clouds in the sky. Sparkling, digital artifacts, symbols and numbers fall from the sky like light rain. There is an empty space in front of him, with nothing to see. He wears a long brown duster coat and a black hat, and the color palette is dark blue, grey, and amber tones. The shot was captured using a Sony A7S III camera with cinematic lighting." I then brought the image into Photoshop and enhanced the color saturation and applied some additional special effects to enhance the digital artifacts in the air. Also discovered the horse's back hoof was facing the wrong direction—so had to correct that in post as well. Total time to produce approx. 90 min.

Designer: Brian Lydon

Source: Midjourney

Image (p.2)

Origin: Authentic

Technique: Image search, followed by photomontage (two images—photographic and photo illustration—combined using Adobe InDesign image blend settings), total time to produce approx. 15 min.

Designer: Brian Lydon

Source: Shutterstock

Images (pp.4, 14, 21)

Origin: Artificial

Technique: Gen AI has long struggled with consistent character generation across iterations—just one of its many frustrating flaws for creatives seeking reliable results and fast development cycles. Our little, frontier robo-buddy character featured throughout the issue was originally conceived by Midjourney for the cover image—but that model failed miserably in every attempt to pull it out of the photo and isolate it on a white background. Even after going to ChatGPT for recommendations on a prompt that might force Midjourney to cooperate, it still couldn't replicate the figure anywhere close to the original. So, it was time to improvise—ChatGPT was fed the image and its own prompt: "Replicate this exact robot, highly accurate to the reference photo, consistent details, sharp mechanical features, clean metallic textures. Pose: Standing at a three-quarters angle to the camera with his arms by his side looking up and to the left. White, seamless background, studio lighting, no shadows, no extra objects, no text, no alterations."—low and behold, we had our first usable image. That was followed by just a few more prompts and ChatGPT was able to recreate multiple versions of our "clanker" with reasonable consistency. Sometimes you just have to catch these models on a good day. Total time to generate approx. 90 minutes.

Designer: Brian Lydon

Source: ChatGPT

Image (p.5)

Origin: Artificial

Technique: Generated with Midjourney with the prompt: "A cinematic shot of a happy child sitting atop a giant globe, suggestive of the Earth, fascinated and delighted by a glowing incandescent light bulb in front of him that he is playing with like a toy. He is surrounded by small pieces of furniture like he is in his childhood bedroom. Looming in the shadows of outer space, behind and larger than the globe itself, is a giant menacing robot, looking down, expressionless. The shot was captured using a Sony A7S III camera with cinematic lighting. Absolutely no text or words are included in the image." After getting a strong base image, additional manual work was necessary in Photoshop to adjust color, shadows and add cosmos into the background. Total time to produce approx. 90 minutes.

Designer: Brian Lydon

Source: Midjourney

Image (p.7)

Origin: Authentic

Technique: Digital photography.

Image (p.9)

Origin: Artificial

Technique: Generated with Midjourney, after uploading a picture of our guest contributor Chris Phenner, with the prompt: “A cinematic shot of a Tron-movie-style setting, featuring a moody and atmospheric futuristic and dark cityscape but with neon highlights. “Light bikes” similar to those featured in Tron race past in the background, causing dramatic, neon streaks of light. In the foreground stands this man, completely clipped out of the reference photo (do not use any other aspects of the reference photo except the man), he is wearing a futuristic outfit, looking intently at his mobile phone as though unaware of his surroundings. The shot was captured using a Sony A7S III camera with cinematic lighting. Absolutely no text or words are included in the image.” Midjourney struggled quite a bit in its attempts to achieve a decent combination of likeness and setting, leading to hundreds of iterations. Total time to generate approx. 90 minutes.

Designer: Brian Lydon

Source: Midjourney

Image (p.12)

Origin: Artificial

Technique: Generated with Midjourney with the prompt: “A candid shot of a female office worker using a computer to enter AI prompts. High resolution (8K or UHD), soft natural lighting and realistic depth of field. The shot was captured using a Sony A7R IV, 50mm f/1.8 lens, ISO 100.” Total time to generate approx. ten min.

Designer: Brian Lydon

Source: Midjourney

Image (p.16)

Origin: Artificial

Technique: Generated with Midjourney with the prompt: “A fantastical, cinematic shot of a fishing boat at sea during a raging tempest. Lightning strikes in the distance. The form of a giant, digital sea serpent is barely visible amongst the waves. The ocean churns around the vessel with incandescent, sparkling digital artifacts, numbers and characters, as though the water is made of computer code. The shot was captured using a Sony A7S III camera with cinematic lighting.” After hundreds of iterations, Midjourney was not able to generate a convincing “digital sea serpent,” so that was passed over in favor of a strong boat and tempest scene—but manual work to combine two images in Photoshop was required in order to get a convincing finished product. Total time to produce approx. two hours.

Designer: Brian Lydon

Source: Midjourney

Images/video clips (p.18)

Origin: Artificial

Technique: Images were created using Sora. Prompts were crafted with the help of a custom GPT within ChatGPT. Some of the prompts were accompanied by a stylistic reference. Once the static images were generated, they were then animated and lip-synched inside of Runway Act-two. The voiceover track was generated with ElevenLabs. Final color grading and editing were completed in After Effects. Total time to generate approx. 2.5 hours.

Designer: Kris Meyers

Source: Sora, Runway Act-two

Video clips (p.18)

Origin: Artificial

Technique: Videos were created using Google VO3. Prompts were crafted with the help of a custom GPT within ChatGPT. Some of the prompts were accompanied by a stylistic reference. Multiple rounds of generations and prompt crafting went into creating the clips. Final color grading and editing were completed in After Effects. Total time to generate approx. 2 hours.

Designer: Kris Meyers

Source: Google VEO3

lydondesign.com