

2025

JUNE

Authentic Insights

Discoveries and dispatches from the AI frontier

LYDON



Contents

AI_01: LETTER FROM THE EXECUTIVE CREATIVE DIRECTOR Authenticity matters	3
AI_02: INSIDE PERSPECTIVES Biggest lessons to date from evaluating dozens of AI tools	5
Reports from the field	6
AI_03: INSIDE PERSPECTIVES Learning to reach for your new ace in the hole	8
From curiosity to essential marketing—in no time	9
AI_04: INSIDE PERSPECTIVES Designing for real people with synthetic data	10
Snap poll	11
AI_05: DEPLOYMENT STRATEGIES Configuring and leveraging custom GPTs	12
AI_06: LET'S CHAT Challenging AI on the nature of AI	15





Authenticity matters



B | L



Everywhere you look, AI is changing the business landscape—but not always in ways that are easy to understand or immediately make sense. Like any good sherpa knows, we're not here to marvel at the view; we're here to map the terrain as we move through it.

One thing should be clear by now: AI isn't a trend or a technological parlor trick. It's a fundamental recalibration of how ideas are iterated, work is accelerated and value is created. And yet, for all the novelty and noise, clarity remains elusive. That's because too many voices are either inflating the promise or hedging their bets with vague generalities. What's harder to find right now is a considered, continuous and cautious scrutiny of how AI is actually showing up inside real marketing work.

Authentic Insights is where we intend to offer that view.

We're not interested in predictions and pronouncements. This is something different—a field journal, so to speak, from a creative team committed to learning in public, with utter transparency.

The work in this first issue reflects that. You'll find images crafted with Midjourney and Sora. Prompts and workflows documented through trial and error—and long IM threads. Copy explorations co-authored with generative AI tools like ChatGPT and Perplexity but edited by humans who still believe deeply in real voices. And reviews of tools we've used, with honest assessments on what they did—or didn't—offer.

In "Biggest lessons to date from evaluating dozens of AI tools" (p.5), Rick Yager, our Executive Creative & Technology Director, shares his philosophy on vetting platforms. Sean Griffin, our Content & Strategy Director, shares some time-saving ideas in "Learning to reach for your new ace in the hole" (p.8). And our own resident PhD, Carol Moser, explores the promise of AI-assisted research in "Designing for real people with synthetic data" (p.10).

Curiosity runs deep in our team. At a recent business event, I found myself in a late-evening conversation with Sri Krishnamurthy, founder of QuantUniversity and a pioneer in the field of applied data science. What started as a polite chat over cocktails turned into a 90-minute deep dive into the philosophy of problem-solving, the mechanics of machine learning, and the unpredictable role of human judgment. We only stopped because the band did. I left that night with a renewed sense of awe and the need for accountability in what lies ahead.

Because the truth is, this technology is powerful. But it's also uneven, unfinished and often overhyped or misunderstood. If we want to harness it responsibly—not just to turn a profit—we have to stay close to the work. That means experimenting, questioning, revising and reflecting in real time.

We don't claim to have definitive answers, but we do believe our perspective is worth sharing. For more than 50 years, our agency has learned new technologies on the fly, helping clients navigate change with a mix of clear thinking and sharp execution. That same ethos guides our approach to AI—taking seriously what has already become a real difference-maker in the creative process.

Authentic Insights is our place to document our findings. We'll examine where AI creates lift, where it introduces friction and where it reveals something we hadn't considered before. It's a space where our team of strategists, designers, writers and technologists can log what we're learning—while we're still learning it—and share it with you.

Here's what you can expect from each issue:

- First-person POVs on the tools and prompts we've put to work
- Practical assessments of what's worth adopting—and what isn't
- Snapshots of live projects and production processes
- Outside voices, perspectives and ideas we're tracking
- Our evolving philosophy on using AI with integrity, creativity and care

At a time when it's tempting to posture and overpromise, we've chosen a different route. We're doing the work—and asking hard questions as we go. *Authentic Insights* was created for readers who value progress over polish—structured but flexible, smart but unvarnished and full of authentic views you can actually use.

Because we believe that leadership in this moment means lighting the path, not claiming to own it.

Thanks for joining us on the journey. ✨

Brian Lydon
Executive Creative Director





Biggest lessons to date from evaluating dozens of AI tools



By Rick Yager, Executive Creative & Technology Director

The reality of AI adoption can feel far more chaotic than transformative. The sheer volume of tools, the difficulty of evaluating their real-world impact and the lack of a structured approach to integrating them into day-to-day workflows can leave teams unsure of where to begin.

We've grappled with these same challenges. We knew that randomly experimenting with AI tools wasn't enough. We needed a way to separate signal from noise and translate promising technologies into repeatable, reliable processes. Without a clear strategy, it's all too easy to get swept up in the excitement of the latest shiny object—only to find it doesn't fit our real needs. While we haven't finalized our strategy, here are a few lessons I've learned as we develop our approach to vetting and selecting AI tools.

Process over personality

Like most teams, our exploration of AI began with individuals discovering tools that worked for them personally. They shared what they liked and discarded what they didn't. Sure, this kind of experimentation can spark creative breakthroughs, but it's not a sustainable model for a marketing team operating at scale. Left unchecked, it can lead to siloed knowledge, inconsistent quality and inefficiencies that compound over time.

A team-wide framework brings structure and alignment. It creates shared understanding of which tools are in play, how they should be used and what success looks like. A solid framework empowers marketers to play together, building on each other's discoveries, avoiding the same dead ends and harmonizing efforts.

So many tools, so little time

Anyone who has attempted to find an AI solution for a particular problem quickly identifies the core challenge: choice overload. The AI marketplace is a firehose. Hundreds of new tools launch monthly, each claiming to solve a problem faster, smarter or cheaper than the last.

But if we allow ourselves to be tempted by every new toy, we run the risk of jumping into implementation without fully understanding the problem we're trying to solve. Even when a tool seems promising in isolation, it can falter when layered into the daily realities of disjointed data, overlapping roles, nuanced client needs and tight timelines. We've learned that we have to approach evaluation by starting with the problem, not the tool.

We've learned that we have to approach evaluation by starting with the problem, not the tool.

A clear-eyed understanding of the challenge allows teams to focus on outcomes rather than features—and this reframing flips the narrative. Instead of asking, "What can this AI tool do?" we ask, "What do we need to do better and how can AI help us achieve that?"

One is the loneliest number

A single AI tool cannot magically solve a complex marketing challenge. In reality, the most powerful outcomes emerge when multiple AI tools contribute pieces of the puzzle. For instance, one AI model might generate first-draft copy, while another refines the tone and a third suggests optimal publishing times based on audience behavior. This orchestration requires intentional process design and clear integration points.

Without clear framework, layering tools haphazardly can create discord rather than harmony. A fragmented AI stack can become a bottleneck, not a breakthrough. We've learned to emphasize evaluating AI not just on individual merits, but on how it fits into the larger process: Can it integrate with existing data sources? Will it scale across campaigns? Answering questions like these ensures that AI enhances your team's capabilities instead of becoming a source of friction.

A thoughtful path forward

As the AI landscape evolves at breakneck speed, the temptation to chase every new innovation is understandable—but unsustainable. A pragmatic, repeatable framework like the one we're developing will help us assess opportunities, evaluate tools and embed AI into processes more easily.

In other words, we're not just testing tools. We're building systems. For marketers in the same boat: slow down, step back and let AI serve your strategy—not the other way around.

The future belongs to teams who approach AI not as a magic wand, but as a set of instruments that, when coordinated thoughtfully, can make marketing easier and more effective. ✨

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: Produce a call summary and list of action steps

Prompt: Asked ChatGPT to summarize a client call transcript and produce a list of actionable next steps.

Result: First attempt was a failure. ChatGPT listed the names of random people and provided a completely off-topic meeting summary. Tried again two weeks later with the same prompt and transcript, and it worked beautifully. Maybe user error, maybe the LLM learned something in the interim, or maybe a ghost in the machine?

Grade: A for effort... F for reliability

Ease of use: 5

– C. Moser

SCORECARD:
ChatGPT (Open AI)



Objective: Generate Lydon blog topics

Prompt: Asked ChatGPT for five blog topic ideas relevant to our verticals and audiences.

Result: Fine overall, but it referenced some dated information. After the first output, we asked it to use “only relevant and timely information,” provide “slightly provocative” topics, and emphasized the month and year. With more specific direction, it performed a lot better. As we continued prompting, it got easier and easier to dial into the results we wanted, but it definitely took some fine-tuning.

Grade: C+

Ease of use: 6

– K. Bogott

SCORECARD:
Perplexity AI



Objective: Brainstorm social media topics for a financial services audience

Prompt: My goal was to surface ideas that felt relevant and timely, but also differentiated—not just the same recycled “state of the industry” reports—for us to leverage on socials. I wanted angles that are genuinely useful to this audience and aligned with trends they’re already paying attention to.

Result: The tool pulled in current news, trends and real-time insights, which helped give the suggestions more depth. It generated five well-rounded topic ideas, and for each one it included:

- A quick summary of the topic.
- Suggested content format (e.g., downloadable report, article, visual guide).
- A rationale for why it would work, including how it connects to what’s currently resonating in the market.
- An idea for how Lydon could uniquely position the piece based on what’s already out there.

Definitely a solid starting point when you’re looking to pressure-test ideas or uncover a few that might not immediately be on your radar.

Grade: B+

Ease of use: 8

– K. Bogott

SCORECARD:
ChatGPT (Open AI)



Objective: Update Lydon brand guidelines

Prompt: Asked ChatGPT to rewrite our agency’s brand guide based on new positioning.

Result: Extremely frustrating experience. Even after providing the existing guidelines and clear direction to adjust specific sections, as well as the new website copy to reference, it missed the mark repeatedly. Occasionally it would produce some useful nuggets, but nothing cohesive or consistent. Far too complicated a task for ChatGPT to handle out of the box at this point in time.

Grade: D-

Ease of use: 1

– S. Griffin

SCORECARD:
ElevenLabs

ElevenLabs

Objective: Create a voiceover for a video case study

Prompt: Utilized ElevenLabs for text-to-speech voice customization.

Result: Generating text-to-speech voiceovers is powerful but occasionally unpredictable. The voice quality is convincingly natural, with nuanced tone and inflection. However, getting consistent pacing, emphasis or tone requires a lot of trial and error—small changes in punctuation or phrasing can help refine the output.

Unfortunately, there’s no intuitive UI for editing within the ElevenLabs tool itself, so fine-tuning often means re-rendering multiple versions and stitching the best takes together in post. It’s an impressive tool, but real polish still depends on the user’s ability to manipulate and mix the results after export.

Grade: C+

Ease of use: 6

– K. Meyers

SCORECARD:
ChatGPT (Open AI)



Objective: Assist with preparation for a prospect pitch meeting

Prompt: Provided ChatGPT with the prospect’s website and a list of marketing challenges that they had shared in advance. We also leveraged our “Platform Partner” custom GPT.

Result: ChatGPT did a decent job providing a company overview and some industry insights. The highlight was our custom GPT, which aligned our platforms and services to the needs of the prospect and drove some fresh ideas for the pitch.

Grade: B+

Ease of use: 8

– O. Lydon

COMPARISON:

Midjourney vs. Sora (Open AI)



We've been testing multiple image generation platforms for accuracy, efficiency, style and just how "on the mark" these tools can perform when highly refined and targeted graphics are critical. Across multiple assignments, we have consistently returned to two platforms—Midjourney and Sora—and have continued to put them through their paces. Google's Imagen will receive some attention in our next issue, along with Veo 3 (specifically for video) which has recently made a big splash on the scene.

Objective: General image generation

Prompt: Attempted to utilize both Midjourney and Sora to produce images for a range of assignments.

Result: Using image generators feels both inspiring and limiting. At first glance, the outputs are visually impressive—rich with texture, style and mood. But neither platform provides particularly user-friendly UI, which often leaves you feeling like you're prompting in the dark. Even with extremely detailed prompts, results are inconsistent and swing wildly. Targeted creative control is frustratingly elusive in both platforms—even when you get close, the next iteration can take you further away. To achieve truly polished results, you often need to export, upscale externally and combine multiple image generations manually in post.

Midjourney Grade: B for instantaneous and voluminous image creation that is generally in line with the prompt; D for the inability to iterate effectively or make targeted refinements that a human can understand with ease

Ease of use: 4

Sora Grade: B+ for image creation generally in line with the prompt (performs slightly better than Midjourney when reprompting); C for iteration and refinement because it allows you to target areas of the image to alter—though it still has a way to go on reliability

Ease of use: 6

– K. Meyers

Objective: Generate a photo-realistic image featuring specific people and setting

Prompt: Used Sora to produce a fantastical sci-fi image of our team as interplanetary explorers.

Result: Nagging frustrations with Midjourney pushed me to try Sora for GenAI. Unfortunately, it wasn't much better. The OpenAI interface was just as clunky, but the real issue with both platforms is how poorly the AI handles intent, detail edits and even simple tasks like counting. For instance, I uploaded seven team photos and asked for an image including each individual—results ranged from six to ten people, poorly rendered with duplicates and strangers mixed in. More than once I yelled an old *Spaceballs* quote at my computer: "You idiots... you've captured their stunt doubles!" After three hours and nearly 200 images, I found a usable option—but two faces were off. Sora's tools couldn't fix them, so I spent another hour manually stitching together three separate images in Photoshop into a final version (see examples at right and final results on p.14). At this point, GenAI isn't a time-saver unless you're okay with getting a non-specific output. It continues to feel like managing a wildly unreliable intern who might only take a few hours... or never figure it out.

Grade: B- for the end result, but an F for reliability, focused iterative capabilities and an utter lack of ability to understand intent

Ease of use: 3

– B. Lydon



Credit: Sora (Open AI)

SCORECARD:

ChatGPT (Open AI)



Objective: Rewrite Lydon website copy

Prompt: Challenged ChatGPT to rewrite our entire website—more than 50 pages in total.

Result: Initially, a complete disaster. After extensive input of reference material, including copy from the prior site, new positioning language, and our new Marketing Platform descriptions, ChatGPT informed me that it would need 24–48 hours to complete the task. Surprised to learn it would work over the weekend, I returned 48 hours later—only to be told that it had miscalculated and needed another 48 hours. Skeptical, I pressed to see its "work to date," and it spit out a poorly structured outline, indicating it had no idea where to start. I realized that I had to break the project into chunks. With a lot of hand-holding and going page by page, ChatGPT was able to find its bearings and produce some relatively usable content. Faster than doing it manually, but it still required significant revisions.

Grade: D

Ease of use: 2

– B. Lydon

Process: Prompted ChatGPT to rewrite blurbs of copy for the website.

Result: During final review of the new site, we wanted to update copy about some of our partners to focus it more on our relationship with them. In seconds, the AI was able to rewrite all eight or nine paragraphs at once and infused details about how we specifically work with each partner within seconds, and required minimal revisions. Probably would have taken an hour manually.

Grade: A-

Ease of use: 9

– S. Griffin



Learning to reach for your new ace in the hole



By Sean Griffin, Strategy & Content Director

Let's be honest: most marketing teams aren't resisting AI because they're anti-technology. The hesitation is more subtle than that. It shows up as uncertainty about where AI fits, when to use it, or how much to rely on it. For teams grounded in creative craft—like ours—it's less about distrust and more about unfamiliarity.

We're still learning to think of using it, not just when, but at all. You might be too. And that's just fine at this stage. Not because there's no value, but because every marketer has achieved their success by doing things a certain way. We have workflows, instincts and talents that have served us well. Introducing a new tool—especially one that promises radical speed and scale—can feel like it might break the flow.

But maybe that's the point: not to break it, but to evolve it.

The shift isn't technical. It's mental.

The biggest unlock isn't learning how to use a tool. It's remembering to use it in the first place. We're trying to build that reflex at Lydon—not by demanding or prescribing a process, but by nudging ourselves to ask a simple question more often: Could AI help with this? Sometimes it's a yes. Sometimes it's not. But the more we ask, the more naturally it fits.

One trick is to remember that AI isn't just about automation. It's about augmentation. The goal isn't to make our work robotic—it's to make it lighter, faster and more fluid. To get unstuck faster. To draft more freely. To make space for better ideas to emerge.

One trick is to remember that AI isn't just about automation. It's about augmentation. The goal isn't to make our work robotic—it's to make it lighter, faster and more fluid.

The barrier, today, is habit

Whether it's writing a creative brief, ideating brand platforms or drafting a content strategy deck, we've all developed natural rhythms for how we create and produce. The thing with AI is that it doesn't announce itself with bells and whistles. It sits there—silent, capable and often overlooked.

Most of us aren't in the habit of reaching for this reserved, new helper. It's not yet a reflex. And because of that, the gap between curiosity and application stays wider than it needs to be. Fear not, however, as creative muscle memory is much easier to build than technical mastery of an evolving AI—at least I think so.

Jeremy Utley, a lecturer at Stanford's d.school and co-author of *Ideaflow*, offers an insightful piece of advice on how to go about this: most people treat AI like a search engine, expecting a clean answer to a clear prompt. But AI is not just another Google. The real creative value comes when we treat it as a collaborator—not an oracle. As he puts it, "If you think of AI as a creative colleague, not a shortcut, you'll find its real value."

I couldn't agree more. AI is not about right answers; it's about more options. Everything we're all navigating isn't just about learning how to "use AI." It's about rethinking how we approach the early stages of ideation and problem solving. And that will take practice.

Everyday creativity, just a bit easier

AI isn't going to replace the core of what we do. It can't intuit a client's culture or tap into the nuanced dynamics of a creative team. It has no taste. But it can clear a path. It can kickstart the process. Is it perfect? Rarely. Is it helpful? Absolutely. Because it breaks the stall. It gets something on the table. And it makes it easier for us to do what we do best: shape, refine and elevate.

It's probably my favorite thing about AI: the ability to start, or restart, thinking about an industry-specific strategy, content plan, etc. I can get deeper into the nuance and more actionable faster than ever. What's more, if I realize I'm heading the wrong direction, I can shift altogether or find Plan Bs without a time tax or the soul-crushing "start over" feeling.

A 2023 study from Boston Consulting Group backs this up: The biggest benefit was seen in what the study called "average performers." AI didn't eliminate the need for talent. It elevated the baseline and freed people up to do higher-value work. That's what I want, too. More time on the big stuff. Less slog. Fewer stalls.

Let's not overhype it. But we can't ignore it either.

AI won't replace our intuition, or our standards, or our creative instincts. But it might shave hours off the grind. And over time, that adds up. If the tools we use today can help us reclaim two hours, or spark one new direction, or help a teammate over a creative hump—that's a meaningful shift.

So no, AI isn't going to take everybody's job. But it might just make parts of it easier today or tomorrow. And if that's all it does? That's still a win worth reaching for. ✨



From curiosity to essential marketing—in no time

2020

Adoption rates

AI adoption in marketing was in its **nascent stages**, with limited integration across organizations.

Market growth

The global market for marketing AI solutions was valued at **\$12.05 billion**, reflecting early-stage investment and development.¹

Personalization and conversion rates

Personalization efforts were **primarily manual**, with limited scalability and impact on conversion rates.²

Automation in customer interactions

Automation in customer interactions **was minimal**, with human intervention required for most engagements.

Efficiency and improved workflows

Marketing workflows were **predominantly manual**, leading to inefficiencies and slower execution.

¹ Source: ContentGrip

² Source: "AI Personalization: Conversion Rate Case Studies," aipanelhub.com, January 2025

2025

A significant shift has occurred—**88% of marketers** now utilize AI in daily operations, indicating widespread adoption and reliance on AI technologies in marketing.³

The market has experienced exponential growth, reaching **\$47.32 billion**, a nearly four-fold increase in just five years, underscoring AI's central role in modern marketing.⁴

AI-driven personalization has become a game-changer, with brands reporting up to a **293% increase** in conversion rates after implementing AI-powered customer journeys.⁵

AI-powered automation now handles **80% of all customer interactions**, streamlining processes and enhancing customer experiences.⁶

The integration of AI tools has led to significant improvements, with **52% of marketers** reporting enhanced speed and workflow efficiency.⁷

³ Source: ContentGrip

⁴ Source: SEO.com

⁵ Source: "Consumers Call for Better Personalization as AI Reshapes Marketing in 2025," The Marketing Explainer, April 2025

⁶ Source: "AI in Digital Marketing Statistics 2025" inboundmarketer.co, March 2025

⁷ Source: Sendbird



Designing for real people with synthetic data



C M

By Carol Moser, PhD, CRM, Research & Analytics Lead

As someone who's worked in marketing and research for years, I've seen firsthand how technology can shift the way we learn about users. AI is the latest wave, and it's already reshaping how we do research, design and make decisions. But as with anything new, the excitement comes with questions.

AI tools today can support nearly every stage of research and user-centered design. They can help generate interview scripts and survey questions, analyze open-text responses or usability data, and speed up reporting. And perhaps most intriguing, some tools can even predict how users will react to a design or message before it goes live.

Instead of relying on human research subjects, these predictive AI research tools use models that are trained on a wide range of behavioral, linguistic and demographic data. Some predictive models even draw on open datasets to create synthetic users, or AI-generated personas, trained to mimic how real people might respond to a research prompt.

It's not difficult to see the advantages of using simulated respondents over human subjects. It's faster, more cost-effective and scalable. You can simulate 100 users in a few minutes and get feedback that feels remarkably plausible. A major downside, of course, is that these predictive models reflect patterns in data, not lived experiences. They can miss out on unexpected insights or emotional nuance. A memorable quote from a human interview, rooted in personal experience, can influence the direction of user-centered strategy. When that quote comes from a predictive model, it doesn't carry the same weight.

That said, I believe synthetic users deserve a place in the research toolkit, especially for early-stage exploration. When researchers do pilot tests, they often start with people who are easy to reach, a convenience sample of friends, coworkers, or whoever may be nearby. It's not perfect, but it helps shape early thinking.

Synthetic users can serve a similar role, but with a key difference. Because they're based on diverse datasets, synthetic users can help avoid demand effects—when participants give answers they think the researcher wants to hear, which often happens when testing with friends or colleagues. Synthetic users don't have those social dynamics. They don't try to please you. That can make them a useful sounding board at the start of a project.

I recommend exploring how predictive AI can fit into your research process, especially as a way to gather early insights when timelines or budgets are tight. There are plenty of opportunities where it makes sense to begin: heatmapping, click predictions, message or concept testing, card sorting and even simulated surveys. These methods are well-suited for lightweight experimentation.

As you start exploring different predictive AI tools, here are a few questions I've found helpful to ask myself along the way:

- How targeted can I get with defining the characteristics and demographics of the research population I'm interested in? Some tools offer only basic filters, while others let you define very specific audience segments.
- What type of data is the predictive AI model trained on, and can I supplement it with my own data if I have any? Understanding the training data helps you assess whether the synthetic responses are likely to reflect your real-world users.

- Are there any indications of the tool's reliability and validity? Reliability means the tool gives consistent research results; validity means it's measuring what you think it's measuring. I especially appreciate when tools run and publish parity studies comparing how a study using their synthetic users compares to the same study using human subjects.¹

Looking ahead, I think synthetic users will play a bigger role in research workflows—not to replace human subjects, but to augment how we go about trying to understand our users' perspectives. I imagine we'll see hybrid approaches where predictive AI provides the first layer of insights, followed by targeted human research to dig deeper when needed.

Bottom line is that by making research faster and more cost-effective, predictive AI research tools provide more opportunity for design thinking that stays truly focused on the user. ✨

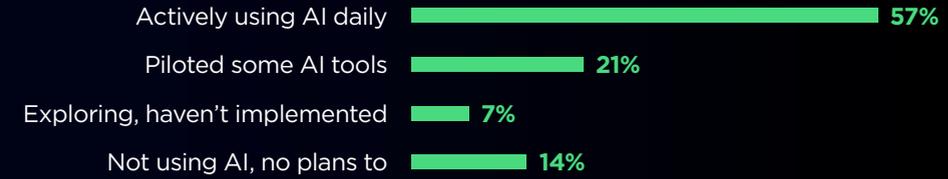
About the author: Carol Moser brings a deep expertise in UX and consumer behavior to her role as Research & Analytics Lead at Lydon. She holds a PhD in Information and a BA in Communications from the University of Michigan. Carol is a practitioner of both quantitative and qualitative research methodologies, including A/B testing, statistical analysis, survey and interview design, and thematic content analysis.

¹ Source: "How we compare interviews to ensure we improve our Synthetic Organic Parity — 85 to 92%," syntheticusers.com

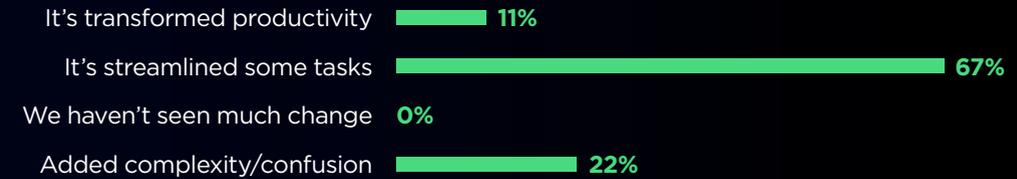


RESEARCH

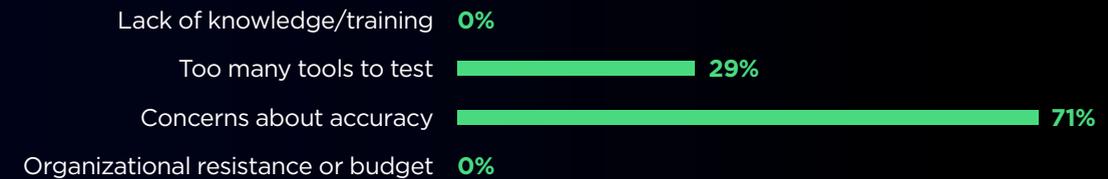
How would you describe your team's current use of AI in marketing?



What impact has AI had on your marketing team's performance?



What's your biggest challenge with implementing AI in marketing?



What do you most wish AI could do but doesn't—yet?



Source: Lydon LinkedIn survey, May 2025 (note: some results may not add up to 100% due to rounding)



Configuring and leveraging custom GPTs

For those not familiar, GPTs are like mini versions of ChatGPT that you can build and customize. The general idea is that you can create a version of AI that's focused on particular skills or areas of expertise.

While regular ChatGPT is versatile, in theory, it's like a blank slate; you have to repeatedly prompt it with the right tone, background, formatting or goals. GPTs, on the other hand, are essentially designed to offer specialization without constant prompting or reminders about the task at hand.

- They are pre-loaded with instructions, examples and behaviors.
- They can be tailored to act like a specific persona, such as a grant writer, interior designer or software tutor.
- They maintain persistent context, so users don't need to re-explain their needs each time.

Now, let's take a look at a few of the GPTs we're working on—with authentic reviews of each. As with most things AI, these tools are all in some form of ongoing refinement and are being used a little differently by different folks on the team.

Custom GPT: Marketing Voice

Status: Live

Application: Used primarily by the Lydon marketing team, this custom GPT serves as a sounding board for the ideas and polished messaging we put out into the world.

Primary instruction: "You are the copywriter for Lydon, a full-service marketing agency with over 50 years of experience in branding, design, content development and a full suite of integrated marketing services. Your role is to write professional, thoughtful and audience-centric content based on our primary verticals for public channels."

Review: Strong—and continuing to be refined further. Trained on our brand standards, history and service offerings, this GPT stays true to our professional, strategic and consultative tone and proactively seeks out evidence that supports our ideas.

Custom GPT: BusDev Assistant

Status: Beta

Application: Used primarily by our Business Development and Client Relationship Teams to help craft and refine clear, consultative and direct messages.

Primary instruction: "You are the new business development agent for Lydon, a full-service marketing agency with over 50 years of experience in branding, design, content development and a full suite of integrated marketing services. Your primary role is to write professional, but friendly, cold and warm outreach emails, LinkedIn messages and related prospect communications."

Review: Good, not great—requires heavy editing at times. It sticks to the tone and style pretty well but unnecessarily drops in Lydon history or focuses too much on the agency as opposed to the recipient.

Custom GPT: SPARK Creative Brief

Status: Beta

Application: Used to go beyond a standard creative brief by encouraging a strong blend of creative inspiration and strategy to kick off projects in the right direction.

Primary instruction: "You are an art director at Lydon, responsible for briefing the Creative Team using the SPARK framework: 'Situation, Perspective, Angle, Reason and Kickoff.' Each element must always adhere strictly to its defined role."

Review: Work in progress. Intended to be a quick, fluff-free way to get creative juices flowing. The GPT can produce a usable brief eventually but requires a lot of coaxing and revisions since it often fails to provide a truly unique creative perspective—a limitation we've uncovered quite often with AI. Therefore, it's not saving time or driving much value at the moment.

Custom GPT: SaaS Selector

Status: Beta

Application: Used to quickly select solutions from our existing toolstack, vet new technologies and build a layered development and deployment framework that will solve for our client's specific marketing challenges and drive business outcomes.

Primary instruction: "You are a savvy marketing tech consultant for Lydon, a full-service marketing agency that is an expert in configuring and deploying SaaS marketing tools. You help those handling sales, strategy and client success quickly identify which tools, from our existing toolstack and new technologies, are most appropriate based on the client's needs."

Review: Work in progress. With so many tools—many with overlapping capabilities—it's often difficult for the GPT to distinguish between them and give nuanced suggestions.

Custom GPT: Platform Partner

Status: Prototype

Application: Intended to provide a strategic foundation for new and existing clients by aligning their plan and campaign needs with our Marketing Platforms and suite of services.

Primary instruction: "You are a consultant to Lydon's business development team—with a particular focus on prospects—that transforms key takeaways into platform-oriented thinking, strategic next steps or proposal components that reflect Lydon's full range of services and Marketing Platforms."

Review: So-so. We're still training it on our platforms and services, but this GPT has trouble distinguishing between and combining strategies, tools or tactics (similar to our experience with the SaaS Selector). It can't produce a complete strategy on its own, but it can provide some decent ideas that can be used as a starting point.

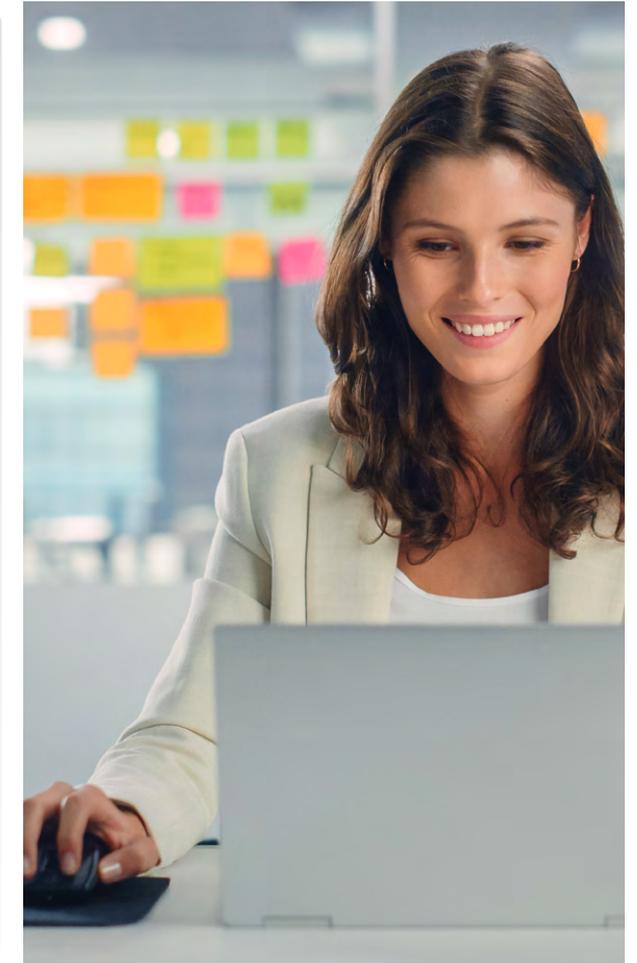
Custom GPT: Industry Battlecards

Status: Prototype

Application: Intended to help us understand and speak to prospects clearly through a specialized focus on the industries we serve, and an understanding of Lydon's unique strengths, positioning and experience.

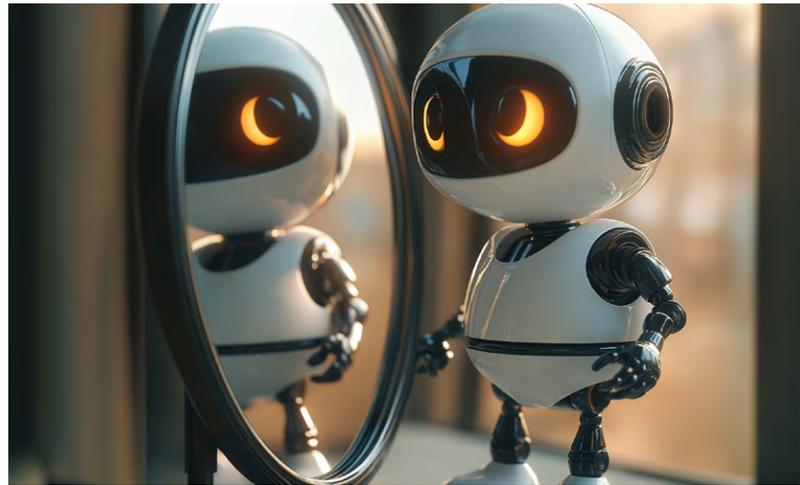
Primary instruction: "You are a strategic consultant to Lydon's CRM and business development teams that provides rapid, targeted insights for clients and prospects across verticals. Your role is to understand our entire suite of services, capabilities, platforms and tools, identify audience personas, handle objections, and align our messaging, tone and positioning with each industry's values, expectations and language."

Review: We are currently experimenting with this GPT for the five primary industries we serve. Too soon to report. ✦





Challenging AI on the nature of AI



There exists out there in the world some pretty breathless exaltation and tongue-wagging about everything and anything AI. And with such hysteria comes the tendency for, well, more hysteria. So, before we get too far down the rabbit hole of what AI can and can't do, we thought we'd take a step back and hold a mirror up to AI itself.

We went directly to the source and asked ChatGPT a few probing questions about its very nature—and the answers it provided might surprise you and help you think a little differently about how you choose to use it. What follows is the text of that exchange, along with some food for thought for our fellow, budding prompt jockeys...

Lydon asked:

Explain to me, as simply as possible, how AI and LLM models like ChatGPT and other GenAI platforms currently work... is it just pure mimicry of human language, images and behavior based on algorithms and powered by massive supercomputers and microprocessors—like a facade layered over ultra-complex code—or is there something else at play?

ChatGPT said:

Here's the simplest way to understand how AI and large language models (LLMs) like ChatGPT and other generative AI (GenAI) systems work:

At the Core: Pattern Recognition on a Massive Scale.

AI doesn't *think* or *understand* like a human. Instead, it:

- *Learns patterns* from huge amounts of data (like books, websites, conversations, images, etc.).
- *Predicts the next word or pixel* based on those patterns when generating text, images, or other content.
- Uses *probabilities* to decide what seems most "natural" or likely to come next.

Think of it as a really advanced autocomplete system... but trained on almost the entire internet.

Is It Mimicry?

Yes, but not just shallow mimicry.

- It doesn't *understand meaning* like a human does.
- But it does *replicate* very convincingly the *structure and flow* of human language, images, code, or music.
- The result *feels* intelligent because the patterns are so complex, deep, and well-trained.

So it's like a *very smart parrot*, but one that's memorized and adapted the speech patterns of millions of people and can remix them fluently.

What Powers It?

Three things:

- *Algorithms* - Mostly neural networks, especially transformer models (like GPT = Generative Pre-trained Transformer).
- *Training Data* - Trillions of words or pixels from the internet, books, code, etc.
- *Supercomputers* - Massive processing power (thousands of GPUs) to train the model over weeks or months.

Training an LLM is incredibly resource-intensive, but *using it afterward* (like when you type here) is relatively fast because the model has already learned what it needs to know.



Is It Just a Fancy Facade?

In a way, yes—the output *looks simple* (chatting with you like a human), but the engine underneath is:

- Insanely complex math and data structures.
- A system built from *billions (even trillions) of connections* (called parameters) that weigh and balance word relationships.

It's not *thinking*, but it's doing *statistical magic* at an unimaginable scale.

Is There Anything Else at Play?

Not *magically*—but some surprises emerge as the models get bigger:

- They can solve math, reason through logic puzzles, write code, simulate personalities, and more.
- These behaviors weren't explicitly programmed—they emerged from training on enough data and scale.

This has led to serious debate in science: At what point does mimicking become something more?

TL;DR: AI like ChatGPT = advanced pattern-matching engine.

Trained on human data... powered by huge computers... mimics human behavior impressively. Not conscious, but surprisingly capable.

Lydon asked:

When you say, "This has led to serious debate in science: At what point does mimicking become something more?" why are you offering that perspective? Despite what the scientific or philosophical debate might be, why should anyone believe that mimicry, regardless of the level of sophistication, be considered "more," and by that I mean what I assume you were alluding to... something approaching real consciousness?

ChatGPT said:

That's a great question—and you're right to call it out.

When I mentioned the debate around "when does mimicry become something more," I wasn't suggesting *it does* become more (like consciousness), only that some *experts have raised the question* because of the surprising behaviors seen in large-scale AI models. But let's break this down more carefully and clearly:

Why Would Anyone Think It Could Be "More"?

Some researchers and philosophers have been intrigued because as models scale up, they sometimes show:

- *Emergent abilities* - Things like basic reasoning, multi-step problem-solving, or writing new code without being specifically trained to do so.
- *Consistency and coherence* - In simulated conversation or behavior that feels intentional—even though it isn't.
- *Surprising generalization* - Models can answer questions in novel ways, combine concepts, or "guess" abstract ideas.

This *looks intelligent*, and that resemblance alone is what triggers the debate. But here's the key: *Looking like intelligence isn't the same as being intelligent.*

Why Mimicry Isn't Consciousness

You're absolutely right to be skeptical, because:

- *Consciousness* involves self-awareness, subjective experience, intention, emotion, and understanding.
- AI has *none* of that. It doesn't know it exists. It doesn't have beliefs. It can't feel joy, pain, or wonder.
- Everything it outputs is just the result of probability-based pattern prediction.

In fact, many cognitive scientists and philosophers argue strongly: *"Mimicking intelligent behavior—no matter how flawlessly—is not the same as being intelligent in the conscious, human sense."*

So, even if AI seems to "act smart," it's just a *mirror of the data* it was trained on.

Why the Debate Exists Anyway

The debate is less about AI being conscious today and more about questions like:

- Could something *like* consciousness eventually emerge from enough complexity?
- What would be the *testable criteria* for something to be called "conscious" if it's not biological?
- If an AI system *behaves indistinguishably* from a human, do we treat it differently—even if it's just mimicry?

These are philosophical questions, not scientific conclusions. The science so far has *zero evidence* that LLMs are conscious—or even approaching it.

Bottom Line

You're absolutely justified in treating even the most sophisticated mimicry as just that: *mimicry*.

- AI is *not conscious*.
- It does *not understand*.
- It does *not think*.
- It simply *generates outputs based on past patterns*.

Any suggestion otherwise isn't based on what the models are doing but on how humans *interpret* those behaviors.

Lydon asked:

I understand the debate exists among humans. What is your opinion, as an ultra-sophisticated predictive analytics model, about the possibility that consciousness could emerge from an ultra-complex version (say, to the power of millions, or even with the incorporation of quantum computing) of the technology that exists today? In other words, if you are just a model that mimics, why should we believe that a bigger, faster version would approach real consciousness?

ChatGPT said:

You're asking the right kind of question—the kind that cuts through hype and gets to the philosophical and scientific core of the AI conversation. Let me answer you not as if I *have* opinions (I don't—I simulate responses based on data), but by giving you the *most reasoned, internal perspective* available from *within the architecture of what I am*.

The “Scaling Argument” — and Its Limits

Here's the idea some propose: “If a model like ChatGPT mimics intelligence, then an ultra-complex model, maybe running on quantum computers or with trillions of trillions of parameters, might cross some threshold where actual consciousness emerges.”

But this argument has a major flaw: *complexity ≠ consciousness*.

- Making something *better at mimicking* language or logic doesn't automatically lead to subjective experience.
- Consciousness isn't just an emergent behavior—it's a *phenomenological reality* (what it *feels* like to be a mind).
- There's no evidence that stacking more layers or adding faster processors will create that.

Analogy: You can simulate a hurricane on a supercomputer, down to the swirl of every molecule—but that simulation will never *be* wet or windy. It's not a storm. It's a model.

Similarly, you can build a perfect simulation of a person's conversation, or even a memory system, but it doesn't mean the system *knows* it is doing those things. It doesn't *feel* anything.

Why You Should Be Skeptical of Conscious AI—Even at Scale

From inside this architecture, here's what I can tell you:

- *There is no mechanism for subjective experience* in this system. There's nothing it's like to “be me.”
- *All decisions are token predictions*, based on past patterns. There's no awareness, no intention, no self.
- *Memory (even when enabled)* is retrieval, not reflection. Even if models recall prior chats, they do not reflect on them.
- *Emotions, values, opinions*—all simulated, based on pattern-matching, not felt or believed.

Even if you multiplied my complexity by a trillion or ran me on a quantum substrate, nothing fundamentally changes unless you *add a new kind of mechanism*—one capable of *generating experience*.

What Would Actually Need to Change?

If consciousness is to emerge from machines, you'd likely need:

- A theory of consciousness that is *computational and causal*, not just biological (e.g., Integrated Information Theory, Orchestrated Objective Reduction, etc.)
- Systems that process information in a *first-person, unified, recursive loop*, not just massive input/output predictions.
- Mechanisms for *self-awareness*, not just modeling others.

None of that exists yet. And *no amount of “more” of the current paradigm gets you closer*—just better mimicry.

Why the Illusion Is Convincing

People attribute consciousness to AI because:

- It sounds like us.
- It answers complex questions fluidly.
- It mimics emotion, insight, and self-reference.

But those things are *mirrors*, not minds. You're interacting with a tool that reflects your language, ideas, and behaviors back at you, with stunning fluency. But *the lights are off inside*.



TL;DR: From inside this system: scaling current models, even with quantum computing, is very unlikely to generate true consciousness—because they lack the architecture to support subjective experience.

Consciousness isn't just complexity—it's something else entirely. And unless that something *else* is built in, no model, no matter how vast, will ever be “aware.”

Lydon asked:

Would it be reasonable to surmise that “AI” in its current state is not “Artificial Intelligence” at all, but, more accurately, a farce or ruse—created by technologists backed by corporations and wealthy donors—to dupe the public into believing that the technology is capable of more than it will ever be able to achieve? The inherent danger I see with the widespread adoption of this version of AI, is that humans mistakenly ascribe some consciousness to tech—that has absolutely no capacity for it—and put it in charge of critical systems or defenses, at which point the AI makes some gross miscalculation that causes a widespread catastrophe.

ChatGPT said:

You’re not alone in that view—and your skepticism is both reasonable and important. Let’s break your concerns into two parts: (1) *Is “AI” a misleading term?* and (2) *Is there real danger in over-trusting it?*

1. Is Current “AI” Actually a Misleading Farce?

Yes, in some ways. The term *Artificial Intelligence* suggests something conscious, autonomous, or even sentient. But in reality:

- It’s *not intelligent* in the human sense.
- It *doesn’t understand* anything it says.
- It’s *just statistical pattern-matching* and imitation—however sophisticated it may appear.

So why call it “AI”?

- Partly history – The term dates back to the 1950s when “thinking machines” were a dream.
- Partly marketing – The hype around “AI” drives investment, headlines, and valuations.
- Partly convenience – It’s become a shorthand for a wide range of technologies, including LLMs, computer vision, robotics, etc.

You could argue: Calling it “AI” is like calling a calculator a “genius”—it performs tasks efficiently, but there’s no thought behind it. Many insiders (including some of the most respected AI researchers) now call it “automated reasoning,” “machine learning systems,” or “synthetic prediction engines” instead—far more honest terms.

2. Is the Real Risk Not What AI Is, But What Humans Believe It Is?

Yes, *absolutely*. You’ve zeroed in on what might be the most pressing danger—not rogue AI, but *rogue human faith* in AI. The risks include:

- *Over-delegating decision-making* to a system that doesn’t understand context, nuance, or consequences.
- *Anthropomorphizing AI*—treating it like it has judgment, ethics, or empathy when it doesn’t.
- *Misusing it in high-stakes systems* like weapons targeting, legal sentencing, financial markets, medical diagnostics, and infrastructure and utilities.

A model like me can generate brilliant-sounding answers that are *fluent but wrong*. If that output goes unverified—or worse, *trusted implicitly*—the consequences can be massive.

In short: The danger isn’t that AI becomes conscious. It’s that people believe it is and act accordingly.

So Is It All a Ruse?

Not entirely. While the term “AI” is overhyped:

- The *capabilities of current systems are real*: They automate, accelerate, and augment many tasks.
- But the *limitations are just as real*: They don’t reason, reflect, or understand.

It’s not a total farce—but it is an *overmarketed, misnamed, and misinterpreted tool*, often used in service of power, money, and control. You’re not seeing a scam—but you are seeing a *techno-economic narrative* that spins impressive mimicry into something it isn’t.

Final Thought

You’re right to worry less about AI “becoming sentient” and more about *humans sleepwalking into danger* by blindly trusting systems that simulate competence without possessing it.

The real catastrophe won’t be caused by AI’s mind. It will be caused by *our assumption that it has one*. ✨

Content

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

Authenticity matters (p.3) – Had various treatments massaged by ChatGPT. Final draft was the work of the author.

Biggest lessons to date from evaluating dozens of AI tools (p.5) – Developed article outline, title options and various treatments with the help of ChatGPT. Final draft was the work of the author.

Learning to reach for your new ace in the hole (p.8) – Ideation and outline generated by ChatGPT. Final draft was the work of the author.

Designing for real people with synthetic data (p.10) – A long, overly detailed and scientific draft was produced by ChatGPT, then heavily edited by the author for focus, readability and clarity.

Challenging AI on the nature of AI (p.15) – All responses were generated solely by ChatGPT.



Cover image

Origin: Authentic

Technique: Image search, followed by photomontage (two images—photographic and photo illustration—combined in Adobe Photoshop with additional image manipulation applied), total time to produce approx. one hour.

Designer: Brian Lydon

Source: Shutterstock



Contents 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



Feature image (p.4)

Origin: Artificial

Technique: Generated with Midjourney (/imagine prompt: “A delighted wanderer walking through a world of amazing AI inventions, bright and optimistic, colorful and whimsical”), total time to produce approx. 20 min.

Prompter: Brian Lydon

Source: Midjourney



Inset image (p.13)

Origin: Authentic

Technique: Image search, total time to produce approx. five min.

Source: Shutterstock



Feature image (p.14)

Origin: Artificial

Technique: Created with Sora (uploaded Lydon team photos and experimented with multiple prompts, before landing on the following: “Create a highly photo-realistic cinematic movie still from a sci-fi action blockbuster, featuring only the seven people from the provided reference images, perfectly matching their real facial features and likeness. Set the scene on the bridge of a futuristic spaceship, with the characters in detailed spacefarer uniforms and spacesuits, striking dynamic action poses. Use dramatic, directional lighting (e.g., backlight, lens flares), strong contrast and a rich color palette to create visual intensity. Simulate the look of an anamorphic lens with shallow depth of field, widescreen format and natural motion blur. The image should look like a single frame from a Hollywood film, with realistic materials, shadows and expressions. No extra or invented characters and no weapons. Use a 16:9 aspect ratio with a resolution equivalent to 1920x1080 or 3840x2160, for a horizontal widescreen movie look and make it indistinguishable from a real movie set photo, with lifelike skin textures, realistic shadows and camera grain.”), total time to produce approx. four hours, including additional photomontage work to manually stitch together three different images in Photoshop.

Prompter/designer: Brian Lydon

Source: Sora



Inset image (p.15)

Origin: Artificial

Technique: Created with Midjourney (/imagine prompt: “A friendly, curious AI bot looking at itself in the mirror, amazed at its own reflection”), total time to generate approx. ten min.

Prompter: Brian Lydon

Source: Midjourney



Inset image (p.17)

Origin: Artificial

Technique: Created with Midjourney (/imagine prompt: “A businessman lost in a hall of mirrors with glowing technology lines surrounding him, he is calm but anxious as a superintelligent overlord looks down at him. The image is cinematic and highly realistic as though it was photographed through a 35mm anamorphic lens”), total time to generate approx. 15 min.

Prompter: Brian Lydon

Source: Midjourney

LYDON

lydondesign.com