Google’s Updates Push Search Further Into Task Completion via @sejournal, @MattGSouthern

Google announced three updates to Search and AI Mode this week, which Roger Montti reported for SEJ. Reading his article motivated me to examine these updates, the broader pattern, and their implications for search this year.

Looking at this in detail, it appears the updates push more of what used to be a results-page experience into task completion.

What Google Announced

Google launched individual hotel price tracking in Search, now available globally for signed-in users searching in English and Spanish. Email alerts notify users of rate changes during selected dates.

Additionally, in March, Canvas trip planning in AI Mode moved from Labs preview to general U.S. availability, allowing users to describe trips and receive custom itineraries with flights, hotels, and attractions that save automatically. Agent-powered store calling, first introduced in classic Search, will soon roll out to AI Mode, enabling Google’s AI to call nearby stores, check inventory, using Gemini models and Duplex.

Rose Yao, Product Leader in Search, posted the updates on X. Additional detail sits in Google’s blog post.

The Pattern

These updates reflect Google's product direction seen in research, patents, and executive statements since January.

In January, Google published the SAGE research paper on training agents for reasoning chains over four steps, laying groundwork for multi-step tasks in Search.

Pichai's April interview made the language public. Pichai said, "A lot of what are just information-seeking queries will be agentic in Search." Our deep dive tracked how his language shifted from "search will change" to specific descriptions of task completion.

Earlier this month, Montti argued that task-based agentic search was already changing SEO, citing Google's global rollout of agentic restaurant booking as evidence that the future tense in Pichai's language was already past tense in product.

A week ago, the U.S. Patent Office published a Google continuation patent titled "Autonomously providing search results post-facto" (our coverage). The filing describes a system that waits for answers when none are immediately available, then delivers them later through assistant interactions.

These updates continue in the same direction. Canvas moves from Labs preview to broader U.S. availability, approximately five months after its initial launch in November. Store calling has been introduced in AI Mode following its debut in Search last November. Additionally, hotel price tracking is now available in Search at the single-property level.

Microsoft's recent news fits the same pattern. Sumit Chauhan, President of Microsoft's Office Product Group, wrote in a company blog post that Copilot's agentic capabilities are now generally available in Word, Excel, and PowerPoint:

> *“Copilot creates the most value when it performs the work—formatting, restructuring, building visuals, and transforming data—rather than just suggesting steps.”*

The features are the default for Microsoft 365 Copilot and Premium subscribers, and available to Personal and Family plans. It’s unclear whether businesses will receive similar reporting for agent-driven surfaces, a point not addressed in Microsoft’s post.

The Vocabulary Hasn’t Settled

Google uses “agentic” in its product language and announcements, describing features like calling and AI Mode as task-oriented. A SeatGeek partnership was called “Google’s Agentic AI Search Experience.” Other companies also use a similar agent framework language.

Pichai describes ‘Agent manager’ as Google’s role and envisions a future in which Search becomes ‘an agent manager’ overseeing various tasks. It positions Google as an orchestration layer on top of agents rather than a direct competitor.

Montti has used “task-based agentic search” in his recent SEJ coverage, sometimes shortened to TBAS. That’s his shorthand for this beat, not industry-standard terminology.

“Agentic” describes the capability. “Agent manager” refers to a specific architectural role that Google is claiming. “Task-based” centers the user’s goal. When three different labels show up in one month, the market is still working out what to call this.

Why This Matters For Search

Professionals

Features introduced this week change the meaning of visibility across several business categories.

Local retailers now encounter a new discovery surface. When a store calls in AI Mode, Google's agents, rather than users, will contact businesses to verify stock and details. Google hasn't disclosed which stores its agents will contact first, how eligibility is decided, or if specific business information influences the process.

An analysis of 68 million AI crawler visits across 858,457 Duda-hosted sites shows that sites with connections to Yext, Google Business Profile, and review systems were crawled more often than those without. These findings describe crawler behavior, not agent calls. It's unknown if similar signals influence which stores are contacted.

Hotels and travel businesses now face individual-property price monitoring. Trip itineraries are based on Canvas's selection logic. No report shows if a hotel appeared in a Canvas plan, triggered an alert, or was named in an AI Mode response.

Publishers face continued pressure from AI-driven summarization. Index Exchange analyzed 1,200 publishers on its exchange platform, finding that 69% experienced year-over-year declines in ad opportunities, with an average drop of 14%.

Declines varied across verticals. Health and careers publishers saw 40-50% ad drops, while news and politics publishers saw only 7% declines.

Vanessa Otero, Founder and CEO of Ad Fontes Media, told Index Exchange for the same piece:

> *"When it's important enough that you want to be accurately and fully informed about some big international, national, or*

> *local event, a quality news site is still a much better experience than asking an AI chatbot, which may give a genericized or inaccurate answer. AI users already know this, which is why most news consumers still go direct to their trusted sites. News has always performed well for advertisers, and if the trend of news site resilience holds, this inventory will likely become the most valuable on the open web of the future."*

Travel publishers face pressure as Canvas compiles itineraries without citing sources, making it impossible for publications to know if their coverage influences trip plans.

Ecommerce retailers lack visibility into which stores get called, so they can't determine if inventory feeds, listing accuracy, or Google Business Profile signals are effective.

Multi-platform coverage complicates strategy. Google's agents favor structured data and verified profiles. Perplexity Computer routes across 19 models with diverse retrieval preferences. ChatGPT Atlas scrapes browser content directly. OpenAI's Operator uses GUI vision to interact with rendered pages.

One business has multiple discovery mechanisms with varying technical needs. Single-strategy optimization no longer covers all surfaces.

What's Still Invisible

Since our coverage flagged the measurement gap, it has widened.

Search professionals still can't see whether their business was included in a Canvas trip plan. They can't see whether an agent called them. They can't see whether their hotel was surfaced in a price-tracking alert. And they can't see how often their content was used to assemble someone else's

itinerary.

No new reporting surfaces were shipped alongside the updates. Alphabet reported $63.1 billion in Google Search & Other advertising revenue for Q4 2025, up 17% year-over-year, with management crediting Search and Cloud acceleration and AI usage gains. No new reporting tools have arrived to help businesses track their role in AI-mediated search.

The pattern holds across platforms. ChatGPT referral data is limited to what OpenAI shares. Perplexity citation visibility is inside Perplexity. Google's agent surfaces don't cleanly map to Search Console.

Academic research on agent training continues to advance. Two April 2026 papers on arXiv show the pace. CW-GRPO, from Junzhe Wang and colleagues, proposes reinforcement-learning improvements for multi-turn search agents. SKILL0, developed by Zhengxi Lu and colleagues at Zhejiang University, trains agents to internalize skill packages. The result is agents that operate without instruction overhead during inference.

The training pipeline is evolving faster than the measurement pipeline businesses depend on. Search professionals can't close that gap alone. Google, OpenAI, Perplexity, and Anthropic would all need to provide equivalent agent-surface reporting. None has publicly committed to doing so.

Looking Ahead

Pichai said that 2027 would be "an important inflection point for certain things." He cited non-engineering workflows and some agentic business processes. Our coverage walked through that timeline.

May brings Google I/O and Microsoft Build. Both companies are likely to expand their agentic surfaces at those events,

making reporting the most urgent thing to watch. If businesses can’t see their role in task-based search, they can’t optimize for it or argue about who should pay for it.

Two longer-running questions sit behind that. Pay-per-click worked when users clicked links. Store calling, Canvas planning, and price tracking don’t produce clicks, and no platform has described a replacement. Schema.org was designed for search engine crawling, not for agents that need real-time inventory, booking availability, and action endpoints. Standards for agent-readable business data haven’t caught up either.

What happens next depends on whether any platform builds the reporting alongside the capability. So far, none has described how it would. Until that changes, businesses will be optimizing for surfaces they can’t see. Next signals land at I/O and Build in three weeks.

More Resources

https://www.searchenginejournal.com/googles-updates-push-search-further-into-task-completion/572888/