

Inside ChatGPT's Confidential Report Visibility Metrics [Part 1] via @sejournal, @VincentTerrasi

A few weeks ago, I was given access to review a confidential OpenAI partner-facing report, the kind of dataset typically made available to a small group of publishers.

For the first time, from the report, we have access to detailed visibility metrics from inside ChatGPT, the kind of data that only a select few OpenAI site partners have ever seen.

This isn't a dramatic "leak," but rather an unusual insight into the inner workings of the platform, which will influence the future of SEO and AI-driven publishing over the next decade.

The consequences of this dataset far outweigh any single controversy: AI visibility is skyrocketing, but AI-driven traffic is evaporating.

This is the clearest signal yet that we are leaving the era of "search engines" and entering the era of "decision engines," where AI agents surface, interpret, and synthesize information without necessarily directing users back to the source.

This forces every publisher, SEO professional, brand, and content strategist to fundamentally reconsider what online visibility really means.

1. What The Report Data Shows: Visibility Without Traffic

The report dataset gives a large media publisher a full month of visibility. With surprising granularity, it breaks down how often a URL is displayed inside ChatGPT, where it appears inside the UI, how often users click on it, how many conversations it impacts, and the surface-level click-through rate (CTR) across different UI placements.

URL Display And User Interaction In ChaGPT

Image from author, November 2025

The dataset's top-performing URL recorded 185,000 distinct conversation impressions, meaning it was shown in that many separate ChatGPT sessions.

Of these impressions, 3,800 were click events, yielding a conversation-level CTR of 2%. However, when counting multiple appearances within conversations, the numbers increase to 518,000 total impressions and 4,400 total clicks, reducing the

overall CTR to 0.80%.

This is an impressive level of exposure. However, it is not an impressive level of traffic.

Most other URLs performed dramatically worse:

- 0.5% CTR (considered “good” in this context).
- 0.1% CTR (typical).
- 0.01% CTR (common).
- 0% CTR (extremely common, especially for niche content).

This is not a one-off anomaly; it’s consistent across the entire dataset and matches external studies, including server log analyses by independent SEOs showing sub-1% CTR from ChatGPT sources.

We have experienced this phenomenon before, but never on this scale. Google’s zero-click era was the precursor. ChatGPT is the acceleration. However, there is a crucial difference: Google’s featured snippets were designed to provide quick answers while still encouraging users to click through for more information. In contrast, ChatGPT’s responses are designed to fully satisfy the user’s intent, rendering clicks unnecessary rather than merely optional.

2. The Surface-Level Paradox: Where OpenAI Shows The Most, Users Click The Least

The report breaks down every interaction into UI “surfaces,” revealing one of the most counterintuitive dynamics in modern search behavior. The response block, where LLMs place 95%+ of their content, generates massive impression volume, often 100 times more than other surfaces. However, CTR hovers between 0.01% and 1.6%, and curiously, the lower the CTR, the better the quality of the answer.

LLM Content Placement And CTR Relationship

Image from author, November 2025

This is the new equivalent of “Position Zero,” except now it’s not just zero-click; it’s zero-intent-to-click. The psychology is different from that of Google. When ChatGPT provides a comprehensive answer, users interpret clicking as expressing doubt about the AI’s accuracy, indicating the need for further information that the AI cannot provide, or engaging in academic verification (a relatively rare occurrence). The AI has already solved its problem.

The sidebar tells a different story. This small area has far fewer impressions, but a consistently strong CTR ranging from 6% to 10% in the dataset. This is higher than Google’s organic positions 4 through 10. Users who click here are often

exploring related content rather than verifying the main answer. The sidebar represents discovery mode rather than verification mode. Users trust the main answer, but are curious about related information.

Citations at the bottom of responses exhibit similar behavior, achieving a CTR of between 6% and 11% when they appear. However, they are only displayed when ChatGPT explicitly cites sources. These attract academically minded users and fact-checkers. Interestingly, the presence of citations does not increase the CTR of the main answer; it may actually decrease it by providing verification without requiring a click.

Search results are rarely triggered and usually only appear when ChatGPT determines that real-time data is needed. They occasionally show CTR spikes of 2.5% to 4%. However, the sample size is currently too small to be significant for most publishers, although these clicks represent the highest intent when they occur.

The paradox is clear: **The more frequently OpenAI displays your content, the fewer clicks it generates.** The less frequently it displays your content, the higher the CTR. This overturns 25 years of SEO logic. In traditional search, high visibility correlates with high traffic. In AI-native search, however, high visibility often correlates with information extraction rather than user referral.

“ChatGPT’s ‘main answer’ is a visibility engine, not a traffic engine.”

3. Why CTR Is Collapsing: ChatGPT Is An Endpoint, Not A Gateway

The comments and reactions on LinkedIn threads analyzing this data were strikingly consistent and insightful. Users don't click because ChatGPT solves their problem for them. Unlike

Google, where the answer is a link, ChatGPT provides the answer directly.

This means:

- Satisfied users don't click (they got what they needed).
- Curious users sometimes click (they want to explore deeper).
- Skeptical users rarely click (they either trust the AI or distrust the entire process).
- Very few users feel the need to leave the interface.

As one senior SEO commented:

"Traffic stopped being the metric to optimize for. We're now optimizing for trust transfer."

Another analyst wrote:

"If ChatGPT cites my brand as the authority, I've already won the user's trust before they even visit my site. The click is just a formality."

This represents a fundamental shift in how humans consume information. In the pre-AI era, the pattern was: "I need to find the answer" → click → read → evaluate → decide. In the AI era, however, it has become: "I need an answer" → "receive" → "trust" → "act", with no click required. AI becomes the trusted intermediary. The source becomes the silent authority.

Shift In Information Consumption

Image from author, November 2025

This marks the beginning of what some are calling “Inception SEO”: optimizing for the answer itself, rather than for click-throughs. The goal is no longer to be findable. The goal is to be the source that the AI trusts and quotes.

4. Authority Over Keywords: The New Logic Of AI Retrieval

Traditional SEO relies on indexation and keyword matching. LLMs, however, operate on entirely different principles. They

rely on internal model knowledge wherever possible, drawing on trained data acquired through crawls, licensed content, and partnerships. They only fetch external data when the model determines that its internal knowledge is insufficient, outdated, or unverified.

When selecting sources, LLMs prioritize domain authority and trust signals, content clarity and structure, entity recognition and knowledge graph alignment, historical accuracy and factual consistency, and recency for time-sensitive queries. They then decide whether to cite at all based on query type and confidence level.

This leads to a profound shift:

- Entity strength becomes more important than keyword coverage.
- Brand authority outweighs traditional link building.
- Consistency and structured content matter more than content volume
- Model trust becomes the single most important ranking factor
- Factual accuracy over long periods builds cumulative advantage

“You’re no longer competing in an index. You’re competing in the model’s confidence graph.”

This has radical implications. The old SEO logic was “Rank for 1,000 keywords → Get traffic from 1,000 search queries.” The new AI logic is “Become the authoritative entity for 10 topics → Become the default source for 10,000 AI-generated answers.”

In this new landscape, a single, highly authoritative domain has the potential to dominate AI citations across an entire topic cluster. “Long-tail SEO” may become less relevant as AI synthesizes answers rather than matching specific keywords. Topic authority becomes more valuable than keyword authority.

Being cited once by ChatGPT can influence millions of downstream answers.

5. The New KPIs: “Share Of Model” And In-Answer Influence

As CTR is declining, brands must embrace metrics that reflect AI-native visibility. The first of these is “share of model presence,” which is how often your brand, entity, or URLs appear in AI-generated answers, regardless of whether they are clicked on or not. This is analogous to “share of voice” in traditional advertising, but instead of measuring presence in paid media, it measures presence in the AI’s reasoning process.

LLM Decision Hierarchy

Image from author, November 2025

How to measure:

- Track branded mentions in AI responses across major platforms (ChatGPT, Claude, Perplexity, Google AI Overviews).
- Monitor entity recognition in AI-generated content.
- Analyze citation frequency in AI responses for your topic area.

LLMs are increasingly producing authoritative statements, such as “According to Publisher X...,” “Experts at Brand Y recommend...,” and “As noted by Industry Leader Z...”

This is the new “brand recall,” except it happens at machine speed and on a massive scale, influencing millions of users

without them ever visiting your website. Being directly recommended by an AI is more powerful than ranking No. 1 on Google, as the AI's endorsement carries algorithmic authority. Users don't see competing sources; the recommendation is contextualized within their specific query, and it occurs at the exact moment of decision-making.

Then, there's contextual presence: being part of the reasoning chain even when not explicitly cited. This is the "dark matter" of AI visibility. Your content may inform the AI's answer without being directly attributed, yet still shape how millions of users understand a topic. When a user asks about the best practices for managing a remote team, for example, the AI might synthesize insights from 50 sources, but only cite three of them explicitly. However, the other 47 sources still influenced the reasoning process. Your authority on this topic has now shaped the answer that millions of users will see.

High-intent queries are another crucial metric. Narrow, bottom-of-funnel prompts still convert, showing a click-through rate (CTR) of between 2.6% and 4%. Such queries usually involve product comparisons, specific instructions requiring visual aids, recent news or events, technical or regulatory specifications requiring primary sources, or academic research requiring citation verification. The strategic implication is clear: Don't abandon click optimization entirely. Instead, identify the 10-20% of queries where clicks still matter and optimize aggressively for those.

Finally, LLMs judge authority based on what might be called "surrounding ecosystem presence" and cross-platform consistency. This means internal consistency across all your pages; schema and structured data that machines can easily parse; knowledge graph alignment through presence in Wikidata, Wikipedia, and industry databases; cross-domain entity coherence, where authoritative third parties reference you consistently; and temporal consistency, where your authority

persists over time.

This holistic entity SEO approach optimizes your entire digital presence as a coherent, trustworthy entity, not individual pages. Traditional SEO metrics cannot capture this shift. Publishers will require new dashboards to track AI citations and mentions, new tools to measure “model share” across LLM platforms, new attribution methodologies in a post-click world, and new frameworks to measure influence without direct traffic.

6. Why We Need An “AI Search Console”

Many SEOs immediately saw the same thing in the dataset:

“This looks like the early blueprint for an OpenAI Search Console.”

Right now, publishers cannot:

- See how many impressions they receive in ChatGPT.
- Measure their inclusion rate across different query types.
- Understand how often their brand is cited vs. merely referenced.
- Identify which UI surfaces they appear in most frequently.
- Correlate ChatGPT visibility with downstream revenue or brand metrics.
- Track entity-level impact across the knowledge graph.
- Measure how often LLMs fetch real-time data from them.
- Understand why they were selected (or not selected) for specific queries.
- Compare their visibility to competitors.

Google had “Not Provided,” hiding keyword data. AI platforms

may give us “Not Even Observable,” hiding the entire decision-making process. This creates several problems. For publishers, it’s impossible to optimize what you can’t measure; there’s no accountability for AI platforms, and asymmetric information advantages emerge. For the ecosystem, it reduces innovation in content strategy, concentrates power in AI platform providers, and makes it harder to identify and correct AI bias or errors.

Based on this leaked dataset and industry needs, an ideal “AI Search Console” would provide core metrics like impression volume by URL, entity, and topic, surface-level breakdowns, click-through rates, and engagement metrics, conversation-level analytics showing unique sessions, and time-series data showing trends. It would show attribution and sourcing details: how often you’re explicitly cited versus implicitly used, which competitors appear alongside you, query categories where you’re most visible, and confidence scores indicating how much the AI trusts your content.

Diagnostic tools would explain why specific URLs were selected or rejected, what content quality signals the AI detected, your entity recognition status, knowledge graph connectivity, and structured data validation. Optimization recommendations would identify gaps in your entity footprint, content areas where authority is weak, opportunities to improve AI visibility, and competitive intelligence.

OpenAI and other AI platforms will eventually need to provide this data for several reasons. Regulatory pressure from the EU AI Act and similar regulations may require algorithmic transparency. Media partnerships will demand visibility metrics as part of licensing deals. Economic sustainability requires feedback loops for a healthy content ecosystem. And competitive advantage means the first platform to offer comprehensive analytics will attract publisher partnerships.

The dataset we’re analyzing may represent the prototype for what will eventually become standard infrastructure.

Image from author, November 2025

7. Industry Impact: Media, Monetization, And Regulation

The comments raised significant concerns and opportunities for the media sector. The contrast between Google's and OpenAI's economic models is stark. Google contributes to media financing through neighbouring rights payments in the EU and other jurisdictions. It still sends meaningful traffic, albeit declining, and has established economic relationships with publishers. Google also participates in advertising ecosystems that fund content creation.

By contrast, OpenAI and similar AI platforms currently only pay select media partners under private agreements, send almost no traffic with a CTR of less than 1%, extract maximum value from content while providing minimal compensation, and create no advertising ecosystem for publishers.

AI Overviews already reduce organic CTR. ChatGPT takes this trend to its logical conclusion by eliminating almost all traffic. This will force a complete restructuring of business models and raise urgent questions: Should AI platforms pay neighbouring rights like search engines do? Will governments impose compensatory frameworks for content use? Will publishers negotiate direct partnerships with LLM providers? Will new licensing ecosystems emerge for training data, inference, and citation? How should content that is viewed but not clicked on be valued?

Several potential economic models are emerging. One model is citation-based compensation, where platforms pay based on how often content is cited or used. This is similar to music streaming royalties, though transparent metrics are required.

Under licensing agreements, publishers would license content directly to AI platforms, with tiered pricing based on authority and freshness. This is already happening with major outlets such as the Associated Press, Axel Springer, and the Financial Times. Hybrid attribution models would combine citation frequency, impressions, and click-throughs, weighted by query value and user intent, in order to create standardized compensation frameworks.

Regulatory mandates could see governments requiring AI platforms to share revenue with content creators, based on precedents in neighbouring rights law. This could potentially include mandatory arbitration mechanisms.

This would be the biggest shift in digital media economics since Google Ads. Platforms that solve this problem fairly

will build sustainable ecosystems. Those that do not will face regulatory intervention and publisher revolts.

8. What Publishers And Brands Must Do Now

Based on the data and expert reactions, an emerging playbook is taking shape. Firstly, publishers must prioritize inclusion over clicks. The real goal is to be part of the solution, not to generate a spike in traffic. This involves creating comprehensive, authoritative content that AI can synthesize, prioritizing clarity and factual accuracy over tricks to boost engagement, structuring content so that key facts can be easily extracted, and establishing topic authority rather than chasing individual keywords.

Strengthening your entity footprint is equally critical. Every brand, author, product, and concept must be machine-readable and consistent. Publishers should ensure their entity exists on Wikidata and Wikipedia, maintain consistent NAP (name, address, phone number) details across all properties, implement comprehensive schema markup, create and maintain knowledge graph entries, build structured product catalogues, and establish clear entity relationships, linking companies to people, products, and topics.

Building trust signals for retrieval is important because LLMs prioritize high-authority, clearly structured, low-ambiguity content. These trust signals include:

- Authorship transparency, with clear author bios, credentials, and expertise.
- Editorial standards, covering fact-checking, corrections policies, and sourcing.
- Domain authority, built through age, backlink profile, and industry recognition.
- Structured data, via schema implementation and rich

snippets.

- Factual consistency, maintaining accuracy over time without contradictions.
- Expert verification, through third-party endorsements and citations.

Publishers should not abandon click optimization entirely. Instead, they should target bottom-funnel prompts that still demonstrate a measurable click-through rate (CTR) of between 2% and 4%, since AI responses are insufficient.

Examples of high-CTR queries:

- “How to configure [specific technical setup]” (requires visuals or code).
- “Compare [Product A] vs [Product B] specs” (requires tables, detailed comparisons).
- “Latest news on [breaking event]” (requires recency).
- “Where to buy [specific product]” (transactional intent).
- “[Company] careers” (requires job portal access).

Strategy: Identify the 10–20% of your topic space where AI cannot fully satisfy user intent, and optimize those pages for clicks.

In terms of content, it is important to lead with the most important information, use clear and definitive language, cite primary sources, avoid ambiguity and hedging unless accuracy requires it, and create content that remains accurate over long timeframes.

Perhaps the most important shift is mental: Stop thinking in terms of traffic and start thinking in terms of influence. Value has shifted from visits to the reasoning process itself. New success metrics should track how often you are cited by AI, the percentage of AI responses in your field that mention you, how your “share of model” compares with that of your competitors, whether you are building cumulative authority

that persists across model updates, and whether AI recognizes you as the definitive source for your core topics.

The strategic focus shifts from “drive 1 million monthly visitors” to “influence 10 million AI-mediated decisions.”

Publishers must also diversify their revenue streams so that they are not dependent on traffic-based monetization. Alternative models include building direct relationships with audiences through email lists, newsletters, and memberships; offering premium content via paywalls, subscriptions, and exclusive access; integrating commerce through affiliate programmes, product sales, and services; forming B2B partnerships to offer white-label content, API access, and data licensing; and negotiating deals with AI platforms for direct compensation for content use.

Publishers that control the relationship with their audience rather than depending on intermediary platforms will thrive.

The Super-Predator Paradox

A fundamental truth about artificial intelligence is often overlooked: these systems do not generate content independently; they rely entirely on the accumulated work of millions of human creators, including journalism, research, technical documentation, and creative writing, which form the foundation upon which every model is built. This dependency is the reason why OpenAI has been pursuing licensing deals with major publishers so aggressively. It is not an act of corporate philanthropy, but an existential necessity. A language model that is only trained on historical data becomes increasingly disconnected from the current reality with each passing day. It is unable to detect breaking news or update its understanding through pure inference. It is also unable to invent ground truth from computational power alone.

This creates what I call the “super-predator paradox”: If

OpenAI succeeds in completely disrupting traditional web traffic, causing publishers to collapse and the flow of new, high-quality content to slow to a trickle, the model's training data will become increasingly stale. Its understanding of current events will degrade, and users will begin to notice that the responses feel outdated and disconnected from reality. In effect, the super-predator will have devoured its ecosystem and will now find itself starving in a content desert of its own creation.

The paradox is inescapable and suggests two very different possible futures. In one, OpenAI continues to treat publishers as obstacles rather than partners. This would lead to the collapse of the content ecosystem and the AI systems that depend on it. In the other, OpenAI shares value with publishers through sustainable compensation models, attribution systems, and partnerships. This would ensure that creators can continue their work. The difference between these futures is not primarily technological; the tools to build sustainable, creator-compensating AI systems largely exist today. Rather, it is a matter of strategic vision and willingness to recognize that, if artificial intelligence is to become the universal interface for human knowledge, it must sustain the world from which it learns rather than cannibalize it for short-term gain. The next decade will be defined not by who builds the most powerful model, but by who builds the most sustainable one by who solves the super-predator paradox before it becomes an extinction event for both the content ecosystem and the AI systems that cannot survive without it.

Note: All data and stats cited above are from the Open AI partner report, unless otherwise indicated.

More Resources:

Featured Image: Nadya_Art/Shutterstock

<https://www.searchenginejournal.com/inside-chatgpts-confidential-report-visibility-metrics-part-1/561608/>