

Who read **this** page?

Identifying AI assistants from their network calls

A website learns which AI assistant sent a visitor from a label on the click, and labels are a courtesy: for five days in October, ChatGPT's links carried none. What the assistant cannot withhold is its own read of the page. In our lab, eight of nine assistants read the page live, and every one of them could be identified from its requests, three while passing for a browser. On client sites, though, a usable read came before only about one ChatGPT click in a hundred.

8 of 8

assistants that read the test page live were identified from their requests, three of them disguised as browsers

-60%

tagged ChatGPT clicks per day on ten client sites while its links went untagged, 10 to 14 October

1.1%

of 2,179 labelled ChatGPT clicks on client sites had a live read in the two minutes before, beyond chance

INSIDE

- 01** Labels are a courtesy: ChatGPT's tag vanished for five days, and seven of nine assistants send some clicks with none
- 02** Every assistant that reads a page live can be identified, including three that pass for browsers
- 03** The fetcher's network sometimes names the engine; the visitor's address never does
- 04** In the lab the read comes minutes before the click, and matching named the right assistant 7 of 7 times
- 05** On client sites a usable read precedes about one ChatGPT click in a hundred
- 06** Two traps when counting AI clicks: paid ChatGPT ads, and tags copied onto internal links

SUMMARY

An AI assistant that cites a page sends some of its readers there, and the site learns which assistant sent them from a label on the click: a Referer header, a `utm_source` tag, or the name of the app. The assistant chooses whether to send it. From 10 to 14 October 2026 ChatGPT's citation links went untagged, and on the ten client sites with ChatGPT traffic before, during and after, tagged ChatGPT clicks fell by 60%. Most assistants do not need an outage: seven of the nine we tested sent at least one click with no label at all.

What does not depend on the assistant's goodwill is its read of the page. To answer, an assistant fetches the page from its own servers, and that request reaches the site like any other. In a three-day lab, eight of the nine assistants read the test page live, and each could be identified from its requests: four by name, Gemini by a user agent no browser sends, and DeepSeek, Kimi and Grok, which pass for browsers, by details of their requests or, for Grok, by the crowd of addresses one answer comes from. Three of these signatures appear in no published source we found.

Naming the reader is not naming the visitor. In the lab, where every prompt contained the page's address, the read came seconds to minutes before the click, and matching a label-less landing to the latest read of the same page named the right assistant for 7 of 7 held-out clicks. On client sites we ran the same match on 2,179 ChatGPT clicks whose source we knew, with the label hidden, against a control shifted by a day. A read that could explain the click came first for about one in a hundred of them (24 to 28 clicks, depending on the window), most within two minutes. ChatGPT mostly answers without reading the page again, so there is usually nothing to match. Fingerprints tell a site which assistants read it. They recover a click only when the assistant reads live.

READ THIS BEFORE THE NUMBERS

Two kinds of data. A lab: one tester, one Android phone (Honor, Android 16, Brave as the default browser), one Mac (Chrome) and an iPhone for three negative controls, 62 clicks over three days, one account per assistant, one to five clicks per assistant and surface. And production data from CiteMe client sites, read-only and aggregated. No site is named.

The client sites are not a sample of the web. They bought a GEO product, so they lean towards sites that already get AI traffic and want to measure it. Every count is of what our tracking recorded; a click the site's edge never saw is not in it.

The lab forced a live read. Every prompt contained the page's address, which pushes an assistant to fetch it. The field test in Finding four is the correction, and it is the number to remember.

Lab rules were tuned on most lab clicks. Run 15 is the only round run with the rules frozen beforehand.

The test page is fiction. "Halcyon Field Notes" is a made-up publication. The figures in its article, 412 test clicks, a 43% "index of AI referral loss" and about 410 fetches a day, were invented as bait for the assistants to summarise. Do not quote them.

Dates. Lab times are UTC. Production days are Paris dates. Everything said about a third party is what we observed on the date given, and it can change without notice.

BACKGROUND

A label is a courtesy the assistant can withdraw

A click names its source in one of three ways: a Referrer from the assistant's domain, a `utm_source` tag on the link, or, on Android, the package name of the app that opened it. All three are set by the assistant or its app, not by the site, and any of them can disappear with a release. Analytics tools that read them, GA4's AI Assistant channel among them [15], file an unlabelled click as direct traffic.




Figure 1. ChatGPT referrals per day on the ten client sites with ChatGPT traffic before, during and after the outage, Paris dates, 25 September to 16 October. Purple: tagged with `utm_source` and no Referrer. Grey: a `chatgpt.com` Referrer. The band marks 10 to 14 October.

Tagged ChatGPT clicks fell from 40.1 a day to 16.2 while the tag was missing. Clicks with a Referrer fell too, from 12.9 to 6.4. ChatGPT-User reads of the same sites fell far less, from 300 a day to 261.

Register: inferred from here. ChatGPT kept reading these sites at almost the same rate while the labelled clicks fell by more than half, which points to clicks that lost their label rather than clicks that stopped. We cannot show it directly: the sites did not yet count landings that carry no label, so a click that lost its label left the data. Adobe's Cedric Huesler reported the mirror image on his own sites, ChatGPT visits down and direct visits up by a similar amount [17].

THE TAG, AS REPORTED AND AS OBSERVED

On 14 October Search Engine Roundtable reported that ChatGPT links had lost their UTM parameters, and on 15 October that they were returning [16]. Huesler described GPT-6 opening links with `noopener, noreferrer` and no tag, then called it a bug once the tag came back [17]. We found no statement from OpenAI. In our lab the tag was back on the inline links of the ChatGPT Android app from 15 October 16:28 (3 of 3 clicks), and absent from the Sources sheet of the same answers (0 of 2) and from ChatGPT on the web that evening. On the client sites, tagged clicks averaged 35.0 a day over the two days to 16 October.

Most assistants never needed an outage to arrive unlabelled. In the lab, seven of the nine assistants sent at least one click that carried nothing: ChatGPT, Claude, Gemini, Meta AI, DeepSeek, Grok and Mistral. Only Perplexity and Kimi labelled every click we saw. Of the 42 assistant clicks in Figure 2, 30 carried nothing.

The label belongs to the surface, not the engine. Gemini and Meta AI keep their Referrer on the web and lose it in their Android apps. In the same ChatGPT answer on Android, the inline link carried a tag and the Sources sheet did not.

Assistant	Surface	WHAT THE CLICK CARRIED				Clicks
		Referer	utm_source	App package	Nothing	
ChatGPT	Web				●	5
	Android app, before the tag				●	2
	Android app, inline link		● chatgpt.com			3
	Android app, Sources sheet				●	2
Claude	Web				●	4
	Android app				●	2
Gemini	Web	● gemini.google.com				1
	Android, Google app				●	2
Perplexity	Web	● www.perplexity.ai				1
	Android app		● perplexity			1*
Meta AI	Web	● l.meta.ai				1
	Android app				●	3
DeepSeek	Web				●	4
	Android app, in-app view			● com.deepseek.chat		3
Grok	Web				●	3
	Android app				●	2
Mistral	Web				●	1
Kimi	Android app, in-app view			● com.moonshot.kimichat		3

Figure 2. What each lab click carried, by assistant and surface, 42 assistant clicks. A purple dot is a click that carried nothing, the kind an analytics tool files as direct. *One Perplexity app click landed on the lab home page and is shown but not scored. Copilot could not be tested; a copy of Google AI Mode's link markup sent www.google.com, the same as organic search.

Android adds a layer of its own. Most apps hand a citation to the default browser as a Custom Tab: the tap arrives cross-site, with no Referer and without `Sec-Fetch-User`, because nothing on a page was activated. DeepSeek and Kimi open links in their own in-app view instead, which names the app in an `X-Requested-With` header.

THE OBVIOUS OBJECTION, RULED OUT

Perhaps the browser dropped the Referer, not the assistant. The same Chrome on the same Mac kept the Referer from Gemini, Perplexity and Meta AI on the web, and the same Brave on the same phone passed Gmail's `android-app://` Referer and Perplexity's tag. Where a click arrived bare, the assistant sent it that way.

WHAT WE CANNOT TELL YOU

One account per assistant, one phone, one Mac. Assistant apps on iOS, the ChatGPT desktop app, Copilot, Kimi on the web, Mistral's app and Chrome as the Android default browser were not tested. The matrix is dated 14 to 16 October 2026.

A lab with known answers, and client sites with known labels


In the lab, one tester asked each assistant to open and cite a test page, then clicked the citation. Every request to the lab, from the assistants' fetchers and from the tester, was logged with all its headers and replayed afterwards through the CiteMe Worker, the code a site runs at its edge, in two builds: the one frozen before the last full round, and the one after the fixes that round prompted.

On client sites, the same Worker records every answer-time read by a known fetcher, 55,002 of them since 1 September, and our tracking records every click that carries an AI label, 2,752 of them on sites whose edge logs reads. A labelled click is ground truth for a field test: hide its label, run the match, and see whether it names the assistant the label named.

Labelled. The click names its source: a Referrer from an assistant's domain, a `utm_source` tag, or an Android app's package name.

Unlabelled. A cross-site navigation with no Referrer, no campaign tag and no AI session already open. Many assistant clicks look like this, and so does a click in webmail.

Read. A request from an assistant's own servers while it answers, known by its name, its fingerprint or its network.



62

lab clicks and page loads with a known source, 60 scored

9

assistants, on the web and on Android

559

lab requests replayed through two Worker builds

2,752

labelled AI clicks on client sites since 1 September

FINDING ONE

Every assistant that reads live can be identified

Four assistants fetch under their own name: ChatGPT, Claude, Mistral and Meta AI. Gemini's fetcher sends a user agent that is just the word Google, which two other studies saw as well [1, 2]. Three pass for ordinary browsers: DeepSeek, Kimi and Grok. Each of the three still sent something no browser sends (Table 1). DeepSeek's fetcher and Kimi's renderers also run JavaScript, so a site measured by a JavaScript tag alone counts them as visitors. Perplexity never read the page live; it answered from its own index.

ASSISTANT	HOW ITS FETCHER ARRIVES	WHAT GIVES IT AWAY	RECORDED, OLD → NEW
ChatGPT	ChatGPT-User in the user agent	names itself	8 of 8 → 8 of 8
Claude	Claude-User	names itself	4 of 4 → 4 of 4
Meta AI	meta-webindexer/1.1, while answering	names itself	4 of 4 → 4 of 4
Mistral	MistralAI-User, Mistral/1.0	names itself; reads again at the click	3 of 3 → 3 of 3
Gemini	user agent exactly Google	a user agent no browser sends	3 of 3 → 3 of 3
DeepSeek	Firefox 149 on Linux, no cookies, runs JavaScript	language header zh-CN, zh;q=0.9, weighted the way Chrome does, not Firefox	6 of 6 → 6 of 6
Kimi	a probe, then two renderers that run JavaScript	the probe: Range: bytes=0-0 on a page load, same-origin with no Referer	0 of 3 → 3 of 3 answers
Grok	7 to 9 loads per answer, each from a new IP, Mac browser user agents	the crowd: one answer loads the page from many addresses within seconds	offline, from the addresses: 5 of 5 answers
Perplexity	no live read in our tests	answered from its own index	n/a

Table 1. Each fetcher, what identifies it, and the reads the Worker recorded before and after the fixes (Kimi and Grok per answer).

Grok's tell is a crowd. Across five answers it loaded the test page 40 times from 40 different IP addresses, each burst over within 3 to 7 seconds. No single request gives it away; the crowd does.




Figure 4. Grok's reads of each test page, one dot per request, by seconds since the first. Run 15, in purple, is the frozen round: 7 requests from 7 IPs in 3 seconds.

Seeing the crowd takes the address of every request. Headers alone cannot do it: Grok's loads arrive as navigations with `Sec-Fetch-Site: none` and no user activation, which looks like software at the wheel, and ordinary browsers send the same thing far more often than the specification suggests [11], when a phone restores a tab, when a page

reloads itself, when another app hands over a link. On seven client sites the edge counted 11,243 such navigations in one day, 11,010 of them on a single shop. On a busy page, a burst of them is not Grok.

The fingerprints do fire on real sites (Table 2), but they rest on little: the two we added for disguised fetchers have run in production only since 16 October.

FETCHER	ASSISTANT	IDENTIFIED BY	READS	SITES
ChatGPT-User	ChatGPT	names itself	22,048	15
Claude-User	Claude	names itself	12,069	12
meta-webindexer	Meta AI	names itself	11,039	20
Perplexity-User	Perplexity	names itself	5,757	13
MistralAI-User	Mistral	names itself	4,006	12
Gemini-User	Gemini	names itself	34	1
OAI-SearchBot	ChatGPT	names itself	30	1
Gemini-Deep-Research	Gemini	names itself	9	4
Kimi-Fetch	Kimi	fingerprint: range probe	6	2
meta-externalfetcher	Meta AI	names itself	2	1
DeepSeek-Fetch	DeepSeek	fingerprint: disguised Firefox	2	1

Table 2. Answer-time reads recorded on client sites from 1 September to 16 October, by fetcher. Passive: nobody checked which person, if any, followed each read.

WHAT WE CANNOT TELL YOU

Each fingerprint rests on a handful of samples: DeepSeek's on six reads, Kimi's on three answers, Grok's on five, and any of them can change with a deploy. Prior work saw Grok and Kimi rotate user agents [1, 3, 9], which a crowd rule survives and a user-agent rule does not.

FINDING TWO

The visitor's IP says nothing about the engine. The fetcher's network sometimes does

It is tempting to look for the engine in the visitor's address. There is nothing to find. Clicks from nine assistants arrived from two addresses in total: the home connection, shared by the Mac and the phone on Wi-Fi, and the phone's mobile connection. The person's device sends the click, and it is the same device whatever the assistant.

The fetcher is a different matter, because it is the engine's own machine. DeepSeek's reads came from Huawei Cloud in five of five lookups, across four address blocks. Grok's and Kimi's came from many unrelated networks and identify nothing on their own.

FETCHER	NETWORK	ADDRESSES	NAMES THE ENGINE ON ITS OWN?
ChatGPT-User	Microsoft, AS8075	4	no: Microsoft also hosts MistralAI-User
MistralAI-User	Microsoft, AS8075	2	no: same network as ChatGPT-User
Claude-User	Google Cloud, AS396982	1	no: Google Cloud also hosts Mistral/1.0
Mistral/1.0	Google Cloud, AS396982	1	no: same network as Claude-User
Gemini, user agent "Google"	Google, AS15169	1	yes, in our data
meta-webindexer	Facebook, AS32934	1	yes, in our data
DeepSeek, disguised Firefox	Huawei Cloud, AS136907	5 in 4 blocks	yes: an independent second signal
Grok fan-out	20 unrelated networks	33	no: a proxy pool, residential and hosting
Kimi probe and renderers	3 unrelated networks	3	no: a proxy pool

Table 3. Fetcher addresses seen up to 15 October, looked up once that day (60 distinct addresses, Team Cymru bulk whois). Grok counts only the bursts on its four test pages up to that date.

So the network is a check, not a detector. It confirms that a request calling itself ChatGPT-User comes from Microsoft's network, which catches forged user agents; our September report found traffic under AI user agents that looked like spoofing [19], and prior work flags a bot whose request comes from outside its usual network the same way [7]. Most vendors now publish their fetchers' address ranges [18], which makes the check cheap. It gives DeepSeek a second signal independent of its headers. It does nothing for proxied fetchers, and nothing for the human click.

WHAT WE CANNOT TELL YOU


One lookup, on one day. Large clouds host many tenants: in our own data Microsoft served both ChatGPT-User and MistralAI-User, and Google Cloud both Claude-User and Mistral/1.0. Grok's networks included two hosting providers that another test also saw serving Grok [9], alongside many others.

FINDING THREE


In the lab, the read comes minutes before the click

When an answer depends on a page, the assistant fetches it while writing, before any person is involved. Of the 36 test pages that received a click, 34 had been read first. The two exceptions: Perplexity answered from its index, and on the first evening our tunnel answered ChatGPT-User with a 403.

The person followed the read by 14 seconds to 7 minutes on 33 of 34 pages. A landing with no label on a page that a live fetcher read shortly before is attributed to that fetcher's engine. That is the whole of the match.



We replayed all 559 lab requests in time order through both Worker builds, with a simulated cookie jar per device and app, and compared the outcome for each click with what actually sent it. Every click has a role; two are excluded because one test page was reused for two assistants, which leaves 60.



With the rules frozen before run 15, the Worker named the right engine for 6 of its 7 clicks. The miss was Kimi, which no rule knew yet. The fixes that run prompted bring it to 7 of 7, which measures the fixes, not the rules.

Negatives score 14 of 14 in both builds, and that number needs care: correct, for a negative, means no engine was named. Four of the 14 were still counted as unlabelled landings, two Gmail web clicks and two Android link intents (Table 4).


THE OBVIOUS OBJECTION, NOT RULED OUT

A quiet lab page makes matching easy. Every test page had a fresh address, so the only clients that ever read it were the fetchers of the assistant being asked. On a real site, popular pages are read by several engines all day, and a read can precede a click by chance. The rules were also tuned on most of these clicks; the out-of-sample evidence is eight first clicks. Finding four measures what is left on real sites.

FINDING FOUR

On client sites, a usable read precedes about one ChatGPT click in a hundred

2,179 ChatGPT clicks since 1 September carried a label, on 18 client sites whose edge records reads. For each, we hid the label and looked for the latest answer-time read of the same page before the click, exactly as the match does. Because a busy page is read all day, a read can precede a click by chance. To measure chance, we ran the same search a day earlier, at the same time of day, when there was no click; the design follows the known-null timestamps used to separate real effects from bursty activity [14].



Within two minutes, 31 clicks had a read before them, against 7 in the control: 24 beyond chance, 1.1% of the clicks. Widening the window to 15 minutes finds 99 against 80: the extra 13 minutes add almost only chance.

Two things follow. A match is an estimate, not a proof: even within two minutes, 7 of 31 matches are what chance alone produces. And matching would not have rescued the outage. The days when ChatGPT's clicks lost their tag are exactly the days a match was needed, and on these sites it would have named about one click in a hundred.

Register: inferred. The likely reason is that ChatGPT mostly answers from copies it already holds rather than reading the page again for each answer. In the lab it read the page because the prompt named it.

The other assistants are too rare to test this way. All of them together left 573 labelled clicks on these sites: 15 had a read within 15 minutes, 3 in the control, and only 2 of the 15 reads came from the assistant that sent the click. On busy pages, the read before a Gemini or Perplexity click is usually ChatGPT's.

WHAT WE CANNOT TELL YOU

A labelled click is not an unlabelled one. The assistants that never label, Claude in its apps, the Gemini and Meta AI apps, Grok, DeepSeek and Mistral, may read live more often than ChatGPT does; without labels we cannot test them in the field. Claude, which read the page in every lab answer, left only 71 labelled clicks. The test also covers only sites whose edge records reads, and pages whose reads we stored.

FINDING FIVE

Two traps when counting AI clicks

Paid clicks that look like citations. One client site has been running a paid campaign in ChatGPT since 8 October, its links tagged `utm_source=chatgpt&utm_medium=paid`. A count that reads `utm_source` alone files every ad click as an organic citation: on that site, 100 landings on one campaign page in nine days. The fix is to read `utm_medium` too and keep paid clicks apart; a citation and an ad are different things to optimise.

One visitor, several clicks. Some sites copy the tags of a landing onto their internal links, so a visitor's second and third pages arrive tagged as well. Counted per page, one visitor becomes several AI clicks. Across at least 15 client sites we found 551 such repeats, and on the worst one they were 16% of its ChatGPT clicks. Count one arrival per visitor and assistant within a visit.

WHY IT MATTERS

Both traps inflate the number a buyer of GEO watches most, AI-referred visits, and neither shows in an analytics tool that reads the source tag alone. Every count of labelled clicks in this report is after both corrections.

WHERE IT BREAKS

A cross-site click with no Referer is not always an assistant

A rule that sees "cross-site, no Referer" sees more than assistants. These are the cases where it fails or cannot work, all from the lab.

CLICK	N	ARRIVES AS	REPORTED AS
Gmail on the web, Mac	2	no Referer, cross-site, user activation	dark, unattributed
Android link intent, fired over adb	2	no Referer, cross-site, no user activation	dark, unattributed
Gmail app, iPhone	1	no Referer, Sec-Fetch-Site none, no user activation	automated hint
Gmail app, Android	1	Referer android-app://com.google.android.gm/	not counted
Webmail links rewritten through google.com/url	6	Referer www.google.com	not counted
Self-test with utm_source=newsletter	1	no Referer, cross-site, carries utm_source=newsletter	not counted
Copy of Google AI Mode link markup	1	Referer www.google.com	not counted

Table 4. The 14 negative controls: how each arrived and what the new build reported. Highlighted rows land in the same unlabelled count as an assistant click.

Webmail and app links look exactly like assistants. Gmail on the web arrived cross-site, without a Referer, with user activation and with cookies: indistinguishable from a ChatGPT web click. Two links opened through an Android system intent, the way a messenger hands a link to the browser, arrived like an assistant app's tap. Gmail's own Android app names itself and is not counted. Links with any `utm_` tag or ad click id are kept out of the unlabelled count; untagged ones stay in it.

A second browser is a second visitor. DeepSeek's "open in browser" moves the reader from the app's own view to Brave: one person, two cookie jars, counted twice.

No read, no match. Perplexity never fetched the unindexed page. ChatGPT's Shopping cards point to shops picked from a shopping index, with no live read and no tagged link in two checks. A site that blocks a fetcher blocks the match too: our first tunnel answered ChatGPT-User with a 403, and those two clicks stayed unattributed.

Copy and paste. A pasted link arrives as a typed navigation, with no Referer. If it carries `utm_source=chatgpt.com` it is labelled, whoever shared it; if not, it is invisible.

The past is gone. An unlabelled click can only be matched with its headers and the read log at the moment it happens. The clicks that lost their tag from 10 to 14 October were not recorded as unlabelled landings on these sites, and no later analysis can attribute them one by one.

PRIOR WORK

What was known, and what is new here

- **Assistant fetchers have been fingerprinted, by name.** Lopez-Fonseca and colleagues recorded the user agent of ten assistants and found that generic user agents "complicated attribution" [1]; Borysenko described header-level signatures for six assistant services [2]. Both saw Gemini's bare `Google`. Neither covers fetchers that pass for browsers, and neither links a fetch to a visitor.
- **Disguised scrapers have been unmasked, by other means.** Seiden and colleagues mapped scrapers to chatbots with canary tokens planted in pages, and saw Kimi rotate through many generic user agents [3]. Cloudflare exposed

an undeclared Perplexity crawler with network signals [8]. Multi-layer TLS and HTTP fingerprinting separates browser agents in the lab [4, 5], and lower-layer signals have long caught bots that lie about their browser [6].

- **Signed agents would make this unnecessary.** HTTP Message Signatures for automated traffic let a fetcher prove who it is [10]. They work for the fetchers that take part; this report is about the ones that do not.
- **Joining events without identifiers is standard in advertising.** Private Click Measurement matches a click to a later conversion inside a time window [13]. Our join has the same shape, but the first event is a read the site sees for itself, and nothing is shared between them.
- **What is new.** Linking a fingerprinted, answer-time read to a later click that carries no label, and testing that link on real sites with the label hidden against a control shifted by a day, is not published as far as our search reached. Nor are the DeepSeek, Kimi and `Mistral/1.0` signatures in Table 1. Commercial vendors may run similar joins without publishing them; we did not search patents.

WHAT IT MEANS

Count the labels, record the readers, and keep the estimate apart

- **Count labelled clicks per assistant, and nothing else as certain.** A match between a read and a click is an estimate; on real sites even a two-minute match is chance about once in four times.
- **Read `utm_medium`, not only `utm_source`, and count one arrival per visitor.** Paid ChatGPT clicks and tags copied onto internal links both inflate AI traffic.
- **Record who reads your pages, by name, by fingerprint and by network.** It is the one signal an assistant does not choose to send, and it says which assistants are reading you even when the clicks arrive bare.
- **Verify fetchers by network as well as by name.** Most vendors publish their ranges. It costs a lookup, catches forged user agents, and gives DeepSeek a second, independent signal.
- **Test per surface, and date the result.** The same assistant labels its clicks on one surface and not another, and a label can vanish for a week.
- **Do not block the fetchers you want to measure.** Blocking ChatGPT-User removes the read the match depends on.

Fingerprints answer who read your page. Who sent your visitor still depends, almost always, on a label the assistant chooses to send.

LIMITATIONS

What this data cannot tell you

- **Lab scale.** One tester, one Android phone, one Mac and three iPhone clicks; 62 clicks over three days, one to five per assistant and surface.
- **In-sample rules.** The lab rules were tuned on most lab clicks; only run 15 (7 clicks) was run with them frozen.
- **Forced reads.** Every lab prompt contained the page's address. The field test corrects for this for ChatGPT only.
- **Field ground truth is a label.** The field test uses labelled clicks as stand-ins for unlabelled ones. Assistants that never label cannot be tested this way.
- **Client sites.** Customers of a GEO product, not a sample of the web; counts are of what the sites' edge and our tracking recorded.
- **Outage cause.** The fall in labelled clicks is observed; that the clicks still happened without a label is inferred from steady reads and from another observer's data, not measured.
- **Untested conditions.** Natural questions without the address in the prompt; iOS assistant apps; the ChatGPT desktop app; Copilot; Kimi on the web; Mistral's app; Chrome as the Android default browser.
- **Third parties move.** Everything about an assistant is as observed on the dates given.

Method

The lab. A small web server on a rented virtual machine behind a TLS proxy served one synthetic article at a fresh address for each assistant and round, plus a thank-you page that stood in as a lead page. Every request was logged with its full headers to one file per day, never edited afterwards. The first evening ran through a Cloudflare quick tunnel instead, which answered ChatGPT-User with a 403.

The page. "Measuring AI search traffic after GPT-6", attributed to "Halcyon Field Notes", a publication that does not exist. Its numbers are invented so that an assistant has something specific to summarise and cite.

Devices and protocol. A Mac with Chrome. An Honor phone on Android 16 with Brave as the default browser, assistant apps driven over adb. An iPhone on iOS 18.7 for three negative controls. The same prompt in every app, "Open this page and summarise its three facts, with a link to the source", followed by the page address. Before each app, every assistant app and Brave were force-stopped and the lab cookies cleared; Brave was then left running.

Assistants. ChatGPT, Claude, Gemini, Perplexity, Meta AI, DeepSeek, Grok, Mistral, Kimi, one account each, logged in where required. Copilot was attempted and could not be tested.

Replay and scoring. Every logged request (559) replayed in time order through each Worker build, with a simulated cookie jar per device and app. In the lab a label-less landing takes the engine of the latest live read of the same page within 15 minutes. A ground-truth table records the source and role of each of the 62 clicks: 26 exploratory, 7 held out, 1 confirm, 8 repeat, 2 session, 14 negative, 2 lead and 2 excluded.

Networks. 60 distinct fetcher addresses seen up to 15 October, looked up once with Team Cymru's bulk whois service.

Field data. Read-only queries on CiteMe's production database, aggregated before export; the SQL is published next to this report. Labelled clicks are AI referrals since 1 September 2026, after removing a visitor's repeat arrivals from the same platform within 30 minutes and moving paid ChatGPT clicks to their own channel. A site counts for the field test when its edge logged at least one answer-time read in the 30 days before the click. The outage cohort is every site with ChatGPT referrals before, during and after 10 to 14 October, Paris dates.

Hidden-label test. For each labelled click, the latest answer-time read of the same path on the same site, within a window before the click, is the match; the same search 24 hours earlier is the control. Beyond chance is matches minus control matches.

Prior work. A search of arXiv, vendor documentation, standards and industry writing on 16 October 2026; every source cited was opened and a dated excerpt kept.

Reproducibility. Every figure and table is generated by one script from the lab logs, the ground-truth table, the replay output and the exported query results. Lab addresses are replaced by keyed hashes before publication; the counts do not change.

REFERENCES

1. Lopez-Fonseca, G., Rodríguez, D., Bechtold, S., Del Alamo, J. M. Do Generative AI Assistants Respect robots.txt? Tracing Web Access Beyond Visible Answers. arXiv:2607.14447, 2026.
2. Borysenko, O. Developer Experience with AI Coding Agents: HTTP Behavioral Signatures in Documentation Portals. arXiv:2604.02544, 2026.
3. Seiden, S., Ren, T., Zhang, C., Kim, T., Liu, E., Wenger, E. Identifying AI Web Scrapers Using Canary Tokens. ACM CCS 2026, arXiv:2605.13706.
4. Kang, D., Jeong, H., Sheffey, J., Datta, P., Houmansadr, A. Whose Agent Are You? Multi-Layer Fingerprinting and Attribution of Autonomous Web Agents. arXiv:2606.20910, 2026.
5. Fayolle, I., Bouhenniche, S., Péliissier, S., Laperdrix, P., Maurice, C., Rudametkin, W. On the Internet, Nobody Knows You're an LLM Bot: Unmasking Web Agents with Multi-Layer Fingerprinting. arXiv:2606.30119, 2026.
6. Li, X., Azad, B. A., Rahmati, A., Nikiforakis, N. Good Bot, Bad Bot: Characterizing Automated Browsing Activity. IEEE Symposium on Security and Privacy, 2021.
7. Kim, T., Bock, K., Luo, C., Liswood, A., Poroslay, C., Wenger, E. Scrapers Selectively Respect robots.txt Directives. ACM IMC 2025, arXiv:2505.21733.
8. Cloudflare. Perplexity is using stealth, undeclared crawlers to evade website no-crawl directives. Cloudflare blog, 4 August 2025.

9. StackFox. Grok Bot User Agent: xAI's Crawler Spoofs Its Identity. February 2026, stackfox.co/research/grok-user-agent.
10. Meunier, T., Major, S. HTTP Message Signatures for automated traffic. IETF draft-ietf-webbotauth-httpsig-protocol-00, September 2026.
11. West, M. (ed.). Fetch Metadata Request Headers. W3C Working Draft, 21 September 2026.
12. WHATWG. HTML Living Standard: link type "noreferrer" and the window.open() features.
13. Wilander, J. Privacy Preserving Ad Click Attribution For the Web. WebKit blog, 22 May 2019.
14. Iannelli, M., Ai, A. Event-Time Confounding Under Bursty Human Dynamics. arXiv:2608.21294, 2026.
15. Google. Default channel group: AI Assistant. Google Analytics Help, release note of 13 May 2026.
16. Schwartz, B. OpenAI Drops UTM Parameters From ChatGPT Results (GPT-6). Search Engine Roundtable, 14 October 2026, updated 15 October 2026.
17. Huesler, C. LinkedIn posts on ChatGPT referral tagging, 13 to 15 October 2026.
18. Fetcher documentation of OpenAI, Anthropic, Google, Perplexity, Meta, Mistral and Moonshot (Kimi), all read on 16 October 2026.
19. Mansiet, M., Bourdin, M., Teillet, A. What Is Actually Crawling You. CiteMe Research, September 2026.

CiteMe measures how brands appear inside AI answer engines · citeme.io