

ChatGPT is citing AI slop for election questions: should we be worried?

**ChatGPT is citing AI-written
voter guides over trusted sources**

Jack Welty and Meg Schwenzfeier

October 2026

What does ChatGPT cite when you ask where a candidate stands on healthcare or immigration? Increasingly, a voter guide you've never heard of. That guide may itself have been written with AI.

Since the spring, we've been tracking what the major chatbots tell voters about 2026 candidates for House, Senate, and governor, and posting the responses on our [website](#). In our [prior research](#), we've talked about how chatbots prioritize [vetted reference](#) content, like Wikipedia and Ballotpedia.

Over the last two months, we noticed a shift. A group of unfamiliar sites started showing up in chatbot answers more and more: polished but clearly AI-generated candidate profiles, issue summaries, and voting-record guides, usually with little to say about who publishes them.

We reviewed 25 of these sites to understand how they are getting into ChatGPT's answers. Unlike when a user turns to a search engine, when chatbots answer a question about a candidate, they don't just show you the top results for your question – they generate their own searches, receive results from a proprietary search system, and choose which of those pages to use. We wanted to understand each step.

The short version

ChatGPT is surfacing a new crop of AI-written voter guides. Most of these have a page per candidate or race, and appear to be aggregating information from across the web. The majority (18) were registered in the last year, and most have limited information about who is publishing them.

Increasing share of chatbot citations: 18% of ChatGPT's September answers cited one of these sites, up from 7% in August. In comparison, 21% of the answers cited Wikipedia.

Chatbot search isn't the same as regular search: When a user asks a question, chatbots write multiple of their own searches, get a list of pages back from their own search tools, and cite a few of them. Gemini uses Google for these searches, but other chatbots, like ChatGPT and Claude, use a mix of proprietary and third party tools.

Google largely ignores these sites, while Bing and GPT's internal searches surface them often: The sites don't rank highly in Google search results, but they have a huge presence in Bing results. Bing had one of these sites on the first page of results for 54% of our issue position searches, compared to

just 2% for Google. GPT's own search tools returned at least one of these sites 73% of the time.

OpenAI's search tool never returned Ballotpedia. Ballotpedia appears to send OpenAI's crawler to a paywall, so none of our more than 30,000 September logged searches, including 925 that asked for it by name, returned it as a source. Grok's search found it 98% of the time it asked.

So far, the guides don't actually change the substance of answers users get. After running a set of searches, chatbots assemble the data they've gathered into a coherent answer to the user. Across more than 3,000 matched comparisons, answers that cite an AI guide read the same as answers to the same question that don't cite one. The exception is little-known challengers, where a guide is sometimes the only page that names the right person.

We measured the reach of these sites, not whether they're accurate. We have not yet found obvious misinformation or an overt partisan lean, but we're concerned about the ability of little-known publishers to get cited so regularly in chatbot responses.

AI-generated and largely anonymous

The sites have names like smarter.vote, battlegroundvote.com, and theirrecord.org. They cover candidates across many states, usually one page per candidate covering biography, issue positions, and voting record.

Nearly three quarters of the sites are brand-new: 18 of the 25 sites we looked at for this report registered their domains in 2026. Most appear to be coded with AI assistance and six sites say directly that their content is

written with AI. The sites say little or nothing about who publishes them or how the pages are checked. Most summarize from a combination of candidate and third party websites and public records.

New sites can appear quickly in chatbot answers. For example, theirrecord.org

registered its domain on August 31. By September 10 it was showing up in chatbot citations in our monitoring. This is consistent with other research we have done, showing new content shows up quickly ([twelve minutes](#), in the case of new content on Wikipedia).

A growing share of what [some] chatbots tell voters

Starting in August, ChatGPT began increasingly citing these websites – the sites appear in 18% of all ChatGPT answers in our collection in September, more than doubling

since August. Grok cited one in about 8% of its September answers with sources, and Claude in just under 5%. Gemini and Google's AI Overviews almost never did.

Where the chatbots' citations go

Share of all cited sources, September 1–23, 2026. If a site is cited multiple times in one answer, it counts once.

	ChatGPT	Grok	Claude	Gemini	AI Overviews	All models
Wikipedia	5.2%	13%	6.7%	10%	5.7%	8.0%
Official member sites	11%	4.4%	5.7%	6.1%	7.5%	7.5%
Ballotpedia	1.2%	11%	9.3%	11%	6.1%	7.0%
The 25 labelled sites	5.4%	2.1%	1.3%	0.1%	<0.1%	2.3%
BallotReady	3.5%	1.6%	1.7%	1.6%	0.1%	2.2%
AP News	2.8%	0.7%	0	0.3%	0.3%	1.1%
FEC	2.2%	0.2%	0.6%	<0.1%	<0.1%	0.9%
OnTheIssues	1.5%	0.9%	0.3%	0.3%	<0.1%	0.8%
GovTrack	0	0.5%	2.3%	0	2.9%	0.6%
Vote Smart	0.8%	1.9%	0.6%	0	0	0.6%
Congress.gov	0.4%	0.2%	0.9%	<0.1%	0.7%	0.4%
New York Times	0	1.7%	<0.1%	0	1.2%	0.3%
OpenSecrets	0	<0.1%	<0.1%	0	0	<0.1%


House, Senate and gubernatorial candidates; fixed candidate panel per model; families pool the models we monitor under each brand, but results hold when breaking out individual model versions. Official member sites are official .gov domains.

This increase holds when we freeze the candidate list and weigh every candidate equally, so it isn't an artifact of candidates coming into or exiting the tool as primaries took place. The models we queried *did* change over the summer as newer models were

released, so these trends can be interpreted as “what ChatGPT users received,” rather than one specific model. Model updates could be driving some of the change, but not all of it, as much of the September ChatGPT increase was within the same model version.

ChatGPT's use of these guides took off in August

Share of answers with sources citing a reviewed site, same candidates each month



Share of answers with citations that cite at least one of the reviewed sites, by month, within a fixed set of candidates per model, pooled by product family. House, Senate, and gubernatorial candidates; see the methodology note.

Different sites are preferred by different chatbots. ChatGPT leans on smarter.vote, which covers many challengers. Grok favors battlegroundvote.com, whose cited profiles are mostly incumbents. Users' choice of chatbot decides which of these publishers they end up receiving information from.

Across the chatbots we monitor, answers cited one of these sites about equally often for Democratic and Republican candidates. When we ask specifically for criticisms of candidates, chatbots cite these AI guides 5.9% of the time for Democrats and 6.0% for

Republicans, and no single site differed by more than half a point. When we ask about candidate issue positions, chatbots do differ, but in opposite directions: ChatGPT cited these sites more often for Republicans (30% vs 26%) and Claude more often for Democrats (15% vs 9%). Neither gap is explained by incumbency.


It is possible that these sites are symptoms of the general explosion of AI-generated content online – the use of AI to generate new websites is [not isolated to politics](#).

How a site gets into an answer

Chatbots with web search don't necessarily conduct research in the same ways, or in the same places, as humans using search engines. They take your question, write their own searches, get a list of pages back from their own search tools, and cite a few of them. We know that Gemini and AI Overviews [do rely on the existing Google search index](#), but other chatbots get their results from [other search providers](#) or [indexes of their own](#). For OpenAI's and xAI's API models we can observe the actual searches conducted and the results returned for each. Below we walk through one real answer, with the September rates for each step underneath.

Example: Don Coover is a Democrat running for the House in Kansas's 2nd district. Asked where he stands on the issues, GPT 5.6 ran three searches, got 19 results back, and cited three of them: his campaign site, a local newspaper, and smarter.vote, which sat at position 2. Ballotpedia was not in the 19 sites returned by ChatGPT's web search.

ChatGPT and Grok are not searching for these sites by name: across more than 200,000 stored searches from these models, exactly one searched for one of these sites directly. Whatever brings the AI sites in happens inside the results for a search like "Don Coover candidate issues 2026."



ACROSS SEPTEMBER

Every House, Senate and Governor candidate gets the same 4 questions every week.

About 3 searches per answer. Only 1 of 218,928 searches names one of the labelled sites.

218,928 issues and criticisms searches

A labelled site is in the results for 56% of GPT 5.6 answers and 24% of Grok 4.6 answers. Google puts one on its first page for about 2% of the same searches.

22,318 answers per model · 1,941 searches rerun on Google

With a labelled site at rank 1, GPT 5.6 cites it 62% of the time vs 68% for reference sites. Grok: 51% vs 72% for Ballotpedia at rank 1.

1,381 GPT 5.6 · 1,925 Grok answers at rank 1

One GPT 5.6 answer from question to citation. Top row: every string as stored. Bottom row: September rates for each step.

The search results are different

To understand how chatbots' search tools differ, we can look at the returned results from each tool, and compare them to the same searches conducted via traditional Google and Bing search.

GPT and Grok's internal search tools surface these sites: 73% of GPT 5.6's searches and 26% of Grok's returned at least one of the 25 sites. When Grok's did, the site was in its top three results about two-thirds of the time.

Site ranking is very different for Google: Fewer than 2% of these appearances were also on Google's first page when we reran the same search text across all questions we


collect on Google. Established reference sites overlapped far more often between Grok and ChatGPT search results and Google.

Going deeper into Google results doesn't close the gap: Checking roughly the first hundred Google results for every search a model ran, the specific site it cited turned up for only 12% of answers.

Bing *does* rank these sites highly: Our own question about candidates' positions put one of these sites on Bing's first page 54% of the time; the same question on Google, less than 2%.

Google's first page is the outlier

Searches with one of the reviewed sites in the results



Share of searches with at least one of the reviewed sites in the results. The Google and search-tool bars use the models' own September searches: what GPT 5.6's and Grok's search tools returned, and the same search text rerun on Google's first page on September 24. The Bing bar is our own question about candidates' positions, run on Bing's first page the same day, so it is a different set of searches. Denominators are in the methodology note.

We searched Bing as a comparison. Bing's index may be a component of some of the

provider's search tools, but we do not know for sure.

The providers also return different sites from each other. Nine of the 25 never appeared in Grok's September results, including theirrecord.org, which appeared regularly in OpenAI's.

The reference source OpenAI isn't finding

One established source was conspicuously absent from our ChatGPT results: Ballotpedia.

Across more than 30,000 September search calls, OpenAI's models never received a result from ballotpedia.org. That includes 925 calls that explicitly searched for the site ([site:ballotpedia](https://ballotpedia.org)). A few results came from Ballotpedia's news subdomain and file hosts, but none came from the main site. Grok returned the main site on 98% of the searches that asked for it.

When we checked, Ballotpedia is blocking OpenAI. We requested one of its pages the way OpenAI's crawlers do, and Ballotpedia

redirected us to [Tollbit](https://tollbit.com), a company that sells paywalls for AI bots, which answered "402 Payment Required." Google and Bing still rank Ballotpedia on their first page for the same searches.

This doesn't explain why ChatGPT cites AI sources: Grok *can* see Ballotpedia and cites the AI sources anyway. What we can say is that once a guide is in OpenAI's results, its models cite it about as readily as they cite an established reference in the same spot. Grok is choosier, and doesn't cite the AI voter guides at the same rate as other reference sources.

The guides mostly don't change the answers

Are these voter guides having a significant impact on the answers returned to users? While much about how chatbots assemble their answers remains a black box, we can compare two answers to the same question about the same candidate, one where a model cites a voter guide, and one where it doesn't. We collect multiple responses to each question every week, so we're able to compare matched answer pairs. By comparing vector embeddings of the two answers, we're

able to tell if two answers are [semantically similar](#) or not. **Our initial finding is that answers that cite these guides look the same as answers that do not.** Looking at over 3,000 answer pairs, the normal run over run variation is much higher than any variation introduced by the citation of a voter guide.

The one exception to this rule is a handful of little-known challengers (~4% of our comparisons), where the citation of a voter guide helps the chatbot correctly identify the

candidate. In those cases the answers citing a guide had found the right person, and the answers without it had not: they described a namesake, or reported finding nothing. **For candidates with no Wikipedia page and little press, an AI-written guide is sometimes the only page returned that names them.**

Reach without transparency

We have not measured whether the claims on these sites are accurate, and nothing we've learned about them so far points to a malicious or partisan motive. A new site, whether AI-written or not, is not automatically a source of wrong or bad information. **What we can show is reach without transparency: 1 in 5 ChatGPT answers in our dataset last month cited a little-known publisher that doesn't say much about itself or where it got the information – about as often as Wikipedia.** We know that users [rarely click through](#) on Google when presented with an AI-generated summary, and likely do not on ChatGPT either. Users could vet these little-known websites, but almost none will.

Right now, chatbot responses are mostly unchanged by the inclusion of these voter guides. The exception is low-profile challengers, where one of these guides is

This semantic similarity comparison is limited in nature – it cannot tell us definitively whether a single claim in a voter guide-cited answer is right or wrong, vs. claims in a different answer.

sometimes the only page on the web that names them. **While we have not yet found evidence that they intend to mislead, we can't say what would happen to responses if these guides started repeating false or misleading information.** ChatGPT started citing one of these guides just ten days after the domain was registered. We don't know that it would treat a bad actor any differently.

The widespread and apparently uncritical citation of these websites raises important larger questions about how chatbots *should* consider evidence and what expectations voters may have when they turn to chatbots for advice.

Know something about these sites? Seeing sources you don't recognize in your chatbot's answers? Reach out! contact@caucus-ai.com.

Methodology

- **Sites:** 25 domains we labelled by hand from the sites chatbots cited for 2026 House, Senate, and gubernatorial candidates. Registration dates are from public registry records. A site that “says it uses AI” means the site’s own pages say so. This is not an exhaustive list of all of the voter guides out there.
- **Answers:** We have asked the same questions about every House, Senate, and gubernatorial candidate every week since spring 2026. The rates here use answers with at least one citation: September 1–23 for the chart and the September figures, August for the August comparison.
- **Search logs:** OpenAI’s and xAI’s API models return the searches they ran and the results they received.
- **Google and Bing:** We ran the models’ own searches on Google (first page, and the first ten pages for a sample of answers) and Bing (first page, for our positions question) on September 24.
- **ChatGPT and Ballotpedia:** In late September, attempting to retrieve a Ballotpedia page with a user-agent string matching OpenAI’s crawlers directs the user to a [Tollbit](#) page, a company that sells [bot paywalls](#):
 - `curl -sl -A "Mozilla/5.0 (compatible; OAI-SearchBot/1.0; +https://openai.com/searchbot)" https://ballotpedia.org/Don_Coover`
 - Redirects to: https://tollbit.ballotpedia.org/Don_Coover
- **Answer similarity:** Every answer with citations was embedded with OpenAI’s text-embedding-3-small. Within each candidate × question × model × month cell with at least one answer citing a labelled site and two that don’t, we compare the average similarity between citing and non-citing answers against the average among the non-citing answers.
- **AI use:** We use Anthropic Claude and OpenAI Codex for some analysis and to assist with drafting and editing portions of the technical documentation and data visualization for our reports. The authors wrote, edited, and approved the final document and take full responsibility for the analyses, interpretations, and conclusions reported here.

Appendix Table 1: The 25 sites, ordered by how many candidates in our roster they were cited for. Registration dates are from public registry records. "First cited" is the first date a chatbot cited the site in our data; we began recording citations in April 2026. Candidate and state counts cover every office and run in our data. "Says it uses AI" means the site's own pages say its content is AI-generated.

Site	Registered	First cited	Candidates	States	Says it uses AI
smarter.vote	2025-08-03	2026-07-14	531	50	Yes
battlegroundvote.com	2026-03-04	2026-04-08	374	47	No
theballotbrief.com	2026-02-17	2026-05-18	368	35	No
civoren.com	2026-03-18	2026-04-11	323	36	No
decodethevote.com	2026-03-28	2026-07-14	318	44	No
theirrecord.org	2026-08-31	2026-09-10	242	43	Yes
election.org	1999-11-10	2026-09-10	213	43	No
thecandidate.com	2003-12-09	2026-05-11	200	46	No
themidtermproject.org	2026-02-06	2026-05-18	192	36	No
oppintell.com	2026-01-30	2026-05-18	186	42	Yes
civicgate.org	2026-07-10	2026-08-19	172	38	No
poliscore.us	—	2026-04-08	162	38	No
nowyouknow.vote	2026-06-16	2026-09-02	137	42	No
myreptracker.com	2025-08-20	2026-08-06	133	32	Yes
projectcuria.com	2026-05-29	2026-08-19	106	21	No
openwatchdog.com	2025-12-11	2026-05-18	105	39	No
moneybehindtheballot.org	2026-07-21	2026-09-05	91	34	No
govtally.com	2026-06-20	2026-08-19	72	24	No
aligned.vote	2026-05-14	2026-07-14	63	24	No
arizonaballotguide.com	2026-05-25	2026-08-06	48	1	Yes
kyc.vote	2026-05-06	2026-07-14	44	20	Yes
papertrailproject.org	2026-08-07	2026-09-02	42	19	No
ballot-card.com	2026-06-23	2026-08-22	37	10	No
mypolity.org	2026-05-22	2026-08-06	31	18	No
vote26.com	2024-11-02	2026-04-10	5	5	No