

Does your site need an llms.txt? An honest read on what it does in 2026

llms.txt is a tidy Markdown map of your site for AI assistants. It is cheap to publish and easy to keep current, but the big answer engines have not committed to reading it yet. Here is what it actually does in 2026, what matters more, and who should bother.

Published 2026-06-30 · A Mue guide

A growing share of buying and research questions now start in an AI assistant rather than a search box, and that has produced a new piece of advice you will see everywhere: add an llms.txt file. The pitch is that it makes your site legible to large language models. Some of that is true, and some of it is hope dressed up as fact. This is what the file actually is, what it does in 2026, and whether it is worth your afternoon.

What llms.txt actually is

llms.txt is a plain Markdown file you serve at the root of your domain, at /llms.txt. It is a curated map of your site for a model with a limited context window: a title, a one-line description, and a short, labelled list of your most important pages, so an assistant does not have to guess which of your URLs matter. The idea was proposed in September 2024 by Jeremy Howard of Answer.AI and is documented at llmstxt.org. A companion file, llms-full.txt, goes further and inlines the actual content of those pages so an assistant can read your material without crawling and rendering each one.

So far so reasonable. The file is small, human-readable, and easy to keep current. The question is not whether it is tidy. It is whether anything reads it.

The honest state of adoption in 2026

Here the marketing and the evidence part ways. The big answer engines have not committed to the convention:

- Google has said its Search systems do not use llms.txt, so it is not a ranking signal and will not change your position in classic results.
- As of 2026, no major AI search provider (Google, OpenAI, Anthropic, or Perplexity) has documented that its crawlers consume llms.txt from external sites.
- Industry monitoring of AI bot traffic has found direct requests to llms.txt files to be a negligible share of all AI crawling, which is the blunt version of the point above.
- Adoption by sites is still early too. One large 2026 crawl of 300,000 domains found only around one in ten had an llms.txt file, after well over a year of industry conversation about it.

None of that makes the file harmful. It makes it a forward-looking bet rather than a present-day ranking lever, and the honest framing is to treat it that way.

Where it genuinely earns its place today

The real traction is not in web search at all. It is in developer tooling, where an assistant is pointed at a documentation site and benefits from a clean index of it:

- Coding agents and IDE assistants increasingly fetch a project or product documentation site on demand, and a tidy `llms.txt` or `llms-full.txt` is exactly the curated entry point they want.
- Several developer-facing companies now publish one for their docs for this reason, which is the clearest signal of where the file is actually useful.
- It doubles as a human-readable index of your best pages, which has modest value regardless of whether any model reads it.

So the pattern is clear: the heavier your documentation and the more developer-shaped your audience, the more an `llms.txt` earns its keep. For a small marketing site the upside is thinner, though the cost is close to zero.

The half that matters right now: can the crawlers even reach you

Before `llms.txt` can help anyone, an AI crawler has to be allowed to fetch your pages at all, and that is decided in `robots.txt`, not in `llms.txt`. This is the part with real consequences today, and it is where sites most often lock themselves out by accident. One distinction changes what you should do:

- Training crawlers such as GPTBot, Google-Extended, and CCBot gather content to train future models. Blocking them keeps your text out of training runs but does not, on its own, remove you from live answers.
- Answer and search crawlers such as OAI-SearchBot, PerplexityBot, and ClaudeBot fetch pages to ground real-time answers and cite sources. If you want to be quoted in AI answers, these are the ones to let in.

The right setting is a deliberate choice, not a default. Plenty of publishers block training while allowing the answer engines. Getting that right is worth more this year than any `llms.txt` file, because it decides whether you can be cited at all.

So should you add one?

Add an `llms.txt` if you have substantial documentation or a developer audience, where coding and doc assistants already read these files. For a small marketing site, add one only because it is cheap and tidy, not because you expect it to move rankings, because it will not. In every case, do the higher-value work first: confirm your `robots.txt` lets the answer engines in, then make your pages easy to parse and quote, which is what genuinely decides whether an assistant cites you. The `llms.txt` is a low-effort bet placed on top of those fundamentals, not a substitute for them.

Frequently asked questions

Do AI assistants actually read `llms.txt`?

Not reliably, not yet. As of 2026 no major AI search provider (Google, OpenAI, Anthropic, or

Perplexity) has documented that its crawlers consume llms.txt from external sites, and industry monitoring of AI bot traffic has found direct hits on llms.txt files to be a negligible share of all AI crawling. Where it is genuinely read is developer tooling: coding agents and in-product assistants that fetch a documentation site on demand. So treat it as a low-cost, forward-looking file, not something that changes how you are answered today.

Is llms.txt the same as robots.txt?

No. robots.txt tells crawlers what they are allowed to fetch; it is an access-control file that has been honoured for decades. llms.txt does the opposite job: it points an assistant that already has access at the pages worth reading and describes them in plain language. robots.txt decides whether the AI crawlers can reach you at all, which is the part that matters right now; llms.txt is a curation layer on top, for the crawlers that choose to read it.

Will an llms.txt improve my Google ranking?

No. Google has said its Search systems do not use llms.txt, so it is not a ranking signal. It will not move your position in classic search results. Its value, if any, is for AI assistants that read it directly, and even there the payoff is unproven. Adding one is reasonable as a cheap, tidy index; expecting it to lift rankings is not.

Who should bother adding an llms.txt?

Sites with substantial documentation or a developer audience get the clearest value, because coding agents and doc assistants do fetch these files, and several developer-tooling companies now publish one for exactly that reason. For a small marketing site the upside is thinner, but the cost is close to zero and it doubles as a clean human-readable index, so it is a fair low-effort bet. Either way, fix your AI crawlability in robots.txt first, since that decides whether you can be cited at all.

Build your llms.txt and check AI crawlers

Free, no signup: agent.mue.app/tools/llms-txt-generator

agent.mue.app/articles/does-your-site-need-an-llms-txt

