# What is llms.txt? An SEO Guide to the AI Web Standard

Published Invalid Date

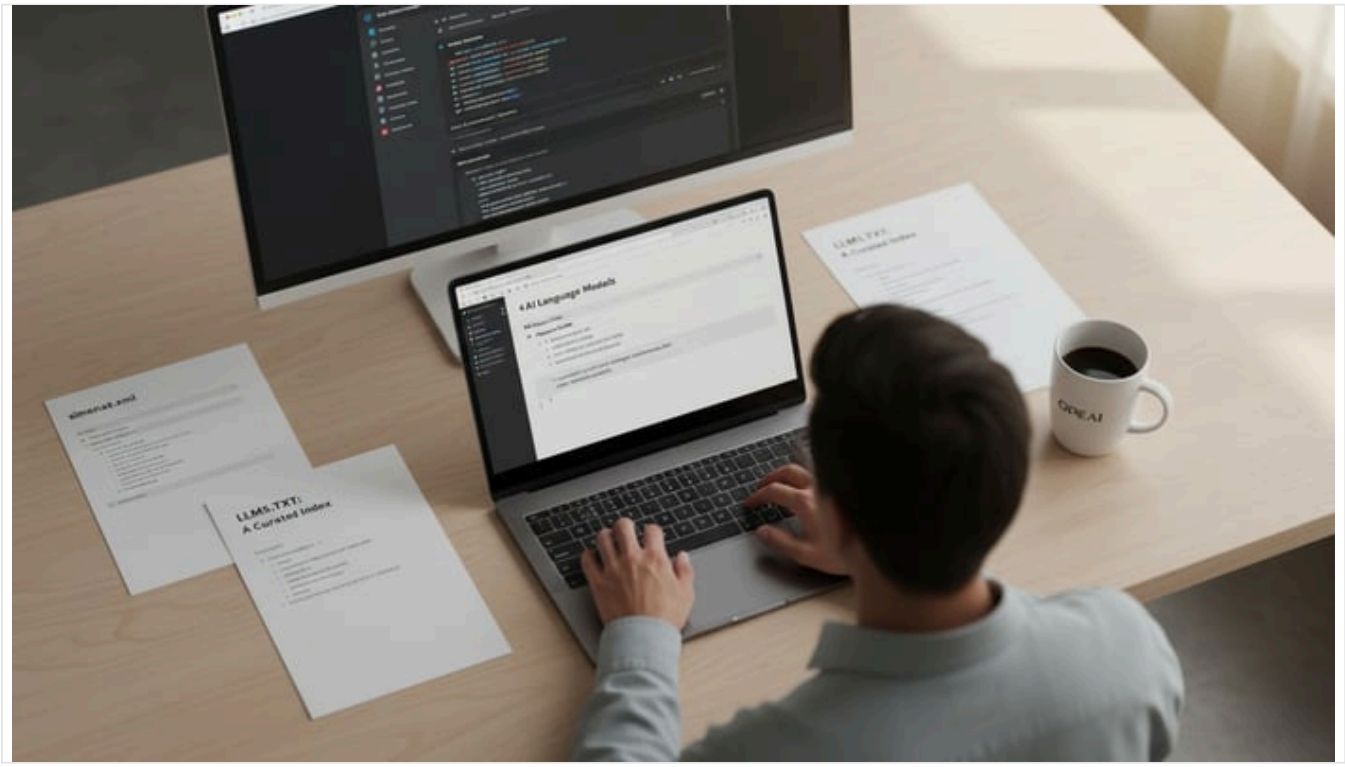


## Executive Summary

The `/llms.txt` file is a newly proposed web standard intended to help *large language models* (LLMs) and AI tools better discover, parse, and interpret website content. Analogous in spirit to the longstanding `robots.txt` for web crawlers, `llms.txt` acts as a **curated, structured map of a site's key pages and information** for AI agents. Proponents argue that, because LLMs have limited context windows and often struggle to extract relevant textual content from complex web pages, a human-authored `llms.txt` can dramatically improve AI accuracy by pointing models directly to the most important, plain-text resources (Source: searchengineland.com) (Source: www.released.so). Early adopters — including developer platforms and some tech companies — have begun creating `llms.txt` files, and tools/generators have emerged to assist implementation (Source: www.released.so) (Source: github.com).

However, **the debate is far from settled**. Some industry voices caution that `llms.txt` may be a premature or unnecessary fix, arguing that traditional search-engine optimization (SEO) already suffices for AI use cases. Google representatives have explicitly stated that Google's AI Overviews rely on standard SEO and will *not* use `llms.txt` (Source: searchengineland.com). Likewise, respected SEO practitioners note that existing mechanisms (e.g. XML sitemaps or creative-commons licenses) can address many needs without a new file format (Source: searchengineland.com) (Source: searchengineland.com). Empirical analysis shows negligible adoption among the top 1,000 websites (effectively 0%) (Source: www.rankability.com) (Source: www.rankability.com), though smaller communities report relatively high "allow AI" policies on sites that *do* implement it (Source: llmscentral.com). Weighing perspectives from AI developers, SEO experts, website operators, and privacy advocates, this report finds that **`/llms.txt` is a compelling innovation in theory but has uncertain practical impact**. Its value will likely depend on whether AI platform maintainers actually heed it, and how web publishers balance the costs of authoring llms metadata against the potential AI outreach benefits.

## Introduction and Background

As generative AI and large language models (LLMs) like OpenAI's GPT and Google's Gemini become pervasive interfaces for information, there is growing interest in making the existing Web more **LLM-friendly**. Currently, websites are primarily built for human readers and traditional search engines; humans easily navigate complex interfaces, and Googlebot indexes pages via links and sitemaps. But LLMs face a critical handicap: *limited context windows*. They cannot ingest entire complex web pages wholesale and often get distracted or confused by navigation bars, ads, scripts, and other non-text elements (Source: searchengineland.com) (Source: llms-txt.io). As Jeremy Howard, the technologist behind the llms.txt proposal, notes:

> *"Large language models increasingly rely on website information, but face a critical limitation: context windows are too small to handle most websites in their entirety. Converting complex HTML pages with navigation, ads, and JavaScript into LLM-friendly plain text is both difficult and imprecise." (Source: searchengineland.com)*

This fundamental limitation means that an AI agent trying to answer a user's question by crawling a site may miss the key information or misinterpret it. Traditional SEO and web design techniques emphasize human usability and search-engine visibility, but they do not directly address the needs of inference-time AI agents (Source: llms-txt.io). In practice, an LLM must sift through page clutter and still can only retain a limited excerpt. For example, one developer reported having to flatten an entire documentation site into a single 115,378-word text file (966 KB) to feed into an LLM with full context (Source: searchengineland.com).

To address this gap, **the `/llms.txt` file was proposed in late 2024** by Jeremy Howard (co-founder of Answer.AI and fast.ai) as a *sympathetic extension* of web metadata standards. The idea is simple: at the root of a website (just as with `robots.txt`), the webmaster can place a plain-text Markdown file named `llms.txt` that contains:

- An H1 title with the site's name or project name (a required element).
- A short introduction or "summary" in blockquote form, giving key context.
- One or more narrative sections to explain the site or usage to an AI.
- Bullet lists under H2 headings, each listing important pages as Markdown links `[Title](URL)` with optional descriptions.
- (Optionally) A separate "Optional" section for lower-priority links the LLM can skip if constrained.

Such a file aims to function as **"a treasure map for AI"** (Source: www.linkedin.com). Instead of forcing AI to parse the website's HTML, the `llms.txt` serves as a **curated table of contents** pointing to all the relevant content. The file itself is written in clear Markdown, stripping out scripts and navigation so the LLM sees only plain text. In practice, an AI agent or tool can fetch "/llms.txt" and see, for example, a title, a summary of the company, then sections like "Products" or "Docs" with bullet links. This gives the model immediate access to the pages and context its creators consider most important.

The notion echoes historical efforts to make the web "understandable to machines." In fact, critics have likened it to the long-dormant **Semantic Web** initiative, which attempted to annotate web content for machine interpretation (Source: news.ycombinator.com). Tim Berners-Lee's decades-old vision of agents "analyzing all data on the Web" in a machine-to-machine "Semantic Web" was never fully realized (Source: news.ycombinator.com). The `llms.txt` approach sidesteps heavyweight ontologies or RDF schemas, instead relying on plain text. As one proponent observed, it avoids the complexity that crushed the Semantic Web effort and uses "stateless formats" (Markdown, XML) to communicate with AI (Source: news.ycombinator.com).

Crucially, `llms.txt` is not about *blocking* or legal control, but about *guiding* AI. Unlike `robots.txt` (which uses "Disallow: URL" rules to prohibit indexing), `llms.txt` has **no blocking directives**. It's entirely optional and instructive – the site owner chooses which pages to highlight. Implementers emphasize it is "rather more of a choosing about which content should be shown contextually or wholly to an AI platform" (Source: searchengineland.com). Effectively, it tells an LLM "if you want to learn our site, here's exactly where to look." For example, Howard and collaborators describe using a small `llms.txt` to feed tools like Cursor or Claude with precisely curated documentation, avoiding the need for each user to manually gather context (Source: news.ycombinator.com).

Thus, `/llms.txt` embodies a **collaborative vision**: websites explicitly collaborating with AI "agents" the same way they collaborate with search engines. As one summary put it, *"LLMs.txt is about to change how your content gets seen, used, and protected in the world of large language models"* (Source: llmsly.com).In this view, it lets content creators "control their narrative" by briefing AI with authoritative info (Source: www.linkedin.com). The proposed benefits range from improved AI answer accuracy to potentially measurable traffic from AI-powered search interfaces. Early experiments by practitioners have given mixed but intriguing signals: engines like OpenAI's models apparently crawl these files, while Google Search (so far) does not automatically use them (Source: searchengineland.com) (Source: searchengineland.com).

However, the `llms.txt` proposal is not universally accepted. Critics point out *elegance vs practicality tensions*. While `llms.txt` may simplify AI crawling, it essentially duplicates what well-designed content should already do: be accessible and clear for *all* readers (human or AI). As one commentator noted, "This isn't good UX for machines. This is a patch for bad UX" – a band-aid rather than fixing the underlying imprecise layouts (Source: news.ycombinator.com). Others worry that without a robust standardization process (e.g. formal registration of a well-known URI or meta tags), the format may splinter. High-profile experts also caution that requiring site owners to hand-author another file burdens them, given that *no AI system currently uses it* (Source: searchengineland.com) (per Google) or has apparently requested such a file. There is even a viewpoint that existing web licensing (Creative Commons etc.) could govern AI use more cleanly than a new text file (Source: searchengineland.com).

In the sections that follow, we delve deeply into **what `/llms.txt` is, how it is supposed to work, and why it may or may not matter**. We examine the technical specification and format (as currently proposed), tools for generation, and differences from related standards like `robots.txt` and `sitemap.xml`. We review the current **state of adoption**, including case studies (e.g. companies trialing `llms.txt` for product docs) and data on how many sites have implemented it. We summarize perspectives from AI developers, SEO specialists, and privacy advocates, using interviews and published statements. We also discuss how AI platforms are reacting – some actively testing `llms.txt`, others remaining agnostic (Source: searchengineland.com) (Source: searchengineland.com). Finally, we lay out potential implications for the future: from how businesses manage their digital content to how search and generative engines will evolve. Through exhaustive citations and analysis, the report aims to answer: *Is `/llms.txt` truly revolutionary for AI search, or just another piece of digital clutter?* Initial evidence suggests it may be important for niche use cases (like developer docs and small sites), but its overall impact in mainstream web discovery remains to be seen.

## The `/llms.txt` Standard: Technical Details and Purpose

The `/llms.txt` proposal and specification are most comprehensively documented by its originators on [llmstxt.org] and associated GitHub repositories (Source: llmstxt.org) (Source: github.com). In brief, an `llms.txt` file is a plain-text **Markdown** document, located at the root of a website (e.g. `https://example.com/llms.txt`). It uses Markdown syntax to present structured content, making it both *human-readable* and parseable by machines. The format intentionally avoids arbitrary nesting or unknown tags, in favor of a well-defined arrangement of headings, paragraphs, blockquotes, and lists. The minimal required element is simply a top-level (H1) heading containing the site or project title (Source: llmstxt.org). Beyond that, the spec defines the following components, in order:

- **H1 Title (required)** – The name of the project or website (e.g. a company name). This anchors the file's identity.
- **Plain-Text Summary (optional)** – A **Markdown blockquote** containing a brief description or vision statement. This "elevator pitch" gives context upfront.
- **Introductory Sections (optional)** – Any number of paragraphs or lists (but *not* additional headings) giving details about the site or instructions for interpreting subsequent links. These can be plain text, bullet lists, etc.
- **H2 Link Sections (optional)** – Zero or more subsections, each headed by an H2. Each H2 is followed by a bullet list of links (Markdown `[text](URL)` anchors), optionally with colon-delimited notes. These compartmentalize the site's content by category. For example:

  ```
  ## Documentation
  - [API Reference](https://example.com/api): Detailed API docs for developers.
  - [Guides](https://example.com/guides): Step-by-step tutorials.
  ```

  Such sections are treated as "file lists" of URLs in the spec; LLMs or tools can iterate through them.
- **Optional "Lower Priority" Section** – It is recommended (but not required) that a final section titled "Optional" list lower-priority pages, so that an LLM can skip them if its context window is limited.

This structure aims to mimic the way humans might summarize a site's information architecture. The file *itself* is written in Markdown specifically *because* Markdown is easily parsed by LLMs and humans alike (Source: llmstxt.org) (Source: golevels.com). The format is unambiguous enough for automated tools to process it using simple text parsing (even regex or XML-based methods as shown by FastHTML's example) (Source: llmstxt.org) (Source: github.com). Critically, the spec emphasizes that the content of `llms.txt` should be concise and relevant — it should not simply dump entire page contents uncritically. Instead, it **highlights** the URLs and facts that the site owner deems most important for AI to ingest.

For instance, the official [llmstxt.org specification] (and [AnswerDotAI's GitHub description]) provides an illustrative mockup:

```
# Example Site Title

> This is a concise summary of the website's purpose and key offerings. It might mention industry, products, or core miss

The following sections list the most important content areas on this site for AI to consider.

## Guides

- [Getting Started](https://example.com/start): An introduction for new users.
- [API Docs](https://example.com/api): The complete API reference.
- [FAQ](https://example.com/faq): Frequently asked questions.

## Projects

- [Project Alpha](https://example.com/alpha): Detailed info on Project Alpha.
- [Project Beta](https://example.com/beta): Overview of Project Beta.

## Optional

- [Blog](https://example.com/blog): News and updates (skip if limited).
```

This example demonstrates the intended use: an AI reading `llms.txt` sees a summary and then clearly structured lists of relevant URLs with short labels or notes. With this, models can pre-load summaries of key pages instead of crawling the entire site blindly.

A key aspect of `llms.txt` is that it **does not attempt to replace web standards**, but to complement them for AI use. For example, it might *implicitly* function like an additional sitemap (listing pages) but with descriptive context. The spec explicitly does *not* define restrictive rules; rather, it is informational. As one explainer notes, `llms.txt` is "similar to robots.txt... but it also offers an additional benefit – full content flattening" (Source: searchengineland.com). In other words, while robots.txt tells machines *what not to crawl*, `llms.txt` tells machines *what to crawl* (and why). It is more akin to an extended **human-curated sitemap** combined with documentation. Indeed, one guidebook formally calls it *"the new robots.txt for the LLM era"* (Source: www.released.so), stressing that it **guides** LLMs to avoid guesswork.

On the practical side, the llms.txt proposal and related tools envision that webpages which have useful content would also offer "clean markdown versions" of those pages (for example at the same URL but with a `.md` extension) (Source: llmstxt.org). This suggestion is like providing pre-processed HTML for machines, but it is not strictly required by the llms.txt standard itself. The primary deliverable of this initiative is the llms.txt file, which may also list optional links (in its sections) to such markdown resources if available. Some projects, like FastHTML, have gone further by programmatically converting their mm-specific pages to Markdown and then referencing them in llms.txt lists (Source: github.com). The FastHTML example is instructive: its `llms.txt` was automatically expanded into "llms-ctx.txt" and "llms-ctx-full.txt" files that incorporate the text from linked pages, tailored for the Claude model's XML context needs (Source: github.com).

In summary, `llms.txt` is a **convention** — not a formal IETF standard (yet) — for how to publish AI-consumable site metadata. It prescribes a specific file name and format, but leaves much flexibility to site owners. The hope is that, by announcing and documenting this convention (via llmstxt.org and GitHub), developers and companies will begin adopting it voluntarily. If enough content providers do so, AI developers (or end-user tools) could programmatically check for `yourwebsite.com/llms.txt` as a known good source of in-page content.

## Relationship to Existing Standards (Robots.txt, Sitemaps, etc.)

To evaluate the significance of `llms.txt`, it is crucial to contrast it with the more established web standards that serve story or search engines. The most natural comparison is `robots.txt`, which has governed web crawler behavior since the 1990s. While both `robots.txt` and `llms.txt` share the idea of a well-known file at the site root, their functions diverge sharply. `robots.txt` is a

**command set** for web robots: it tells search engines (via directives like `User-agent` and `Disallow`) which parts of the site may *not* be scraped or indexed. In contrast, `llms.txt` is **not about blocking**. It provides positive guidance — essentially a quick table of contents — for what *to include* in an LLM's context. As Search Engine Land explains, "robots.txt files work fine for crawlers and do not need changing for the purpose of LLMs" (Source: searchengineland.com), because robots.txt's use case (governing crawl allowances) is orthogonal to that of llms.txt (improving content ingestion).

Another useful analogue is the XML sitemap (`sitemap.xml`). A sitemap is just a list of URLs formatted in XML, optionally with metadata like last-modified dates or priorities, intended entirely for search engines. It does not contain descriptive context or summaries; it simply enumerates pages for discovery. By contrast, `llms.txt` is like a **contextual sitemap**. It still lists links, but in an annotated, human-readable form. A marketer's guide notes that "unlike a sitemap.xml (which is just a list of URLs), llms.txt provides context and structure for each link" (Source: golevels.com). In a way, one can view `llms.txt` as merging the concepts of a sitemap and some form of "About" page: it both enumerates key pages and explains what they are.

We can summarize some key distinctions in the table below:

| ASPECT / FILE | ROBOTS.TXT | SITEMAP.XML | LLMS.TXT |
|---|---|---|---|
| **Purpose** | Control crawler indexing (disallow pages) (Source: searchengineland.com) | Inform search bots of all site URLs and metadata | Curated guide to important content for LLMs (Source: llms-txt.io) (Source: golevels.com) |
| **Content Type** | Plain text directives (e.g. `Disallow:`) | XML with `<url>` entries | Markdown: headings, lists, links, text |
| **Audience/Agent** | Search engine crawlers (Googlebot, etc.) | Search engine crawlers | AI systems and LLM-based agents |
| **Key Difference** | Tells bots *what to skip* | Lists *all pages to include* | Highlights *what to focus on* |
| **Human-readable?** | Yes (simple commands) | No (machine XML format) | Yes (plain Markdown with descriptions) (Source: golevels.com) |
| **Example Use** | `Disallow: /private/` blocks path | `<loc>https://example.com/page.html</loc>` | `- [FAQ](https://exa.com/faq): common topics` |

*(Sources: Consultation of llms.txt proposals and SEO guides (Source: searchengineland.com) (Source: golevels.com) (Source: llms-txt.io).)*

The above highlights that existing standards serve different needs. Traditional SEO optimization (via proper HTML, meta tags, structured data, sitemaps, etc.) remains fundamentally about human users and Google's algorithms (Source: llms-txt.io) (Source: llms-txt.io). `llms.txt` explicitly acknowledges that *those methods are insufficient for AI*. Indeed, as one analysis notes, LLMs "have finite capacity for processing information at once" and "keyword-optimized content doesn't always provide the full understanding LLMs need" (Source: llms-txt.io). In other words, a heavily SEO-optimized site might rank well on Google but still puzzle an AI into missing context or ingesting junk. `llms.txt` is offered as a supplement—*not a replacement*—for SEO practices (Source: llms-txt.io) (Source: llms-txt.io). Good SEO (fast pages, clear headings, etc.) is still necessary for general visibility, while `llms.txt` would additionally ensure *AI* sees the essence of your content.

Other related ideas in the industry support this division. For example, some have suggested adding special `<meta name="LLM">` tags or HTTP header hints to indicate AI-friendly content. One SEO strategist even proposed a `rel="llm"` link or a MIME profile for LLM-friendly markdown (Source: news.ycombinator.com). These proposals share the goal of signaling relevant content to AI, but they

differ in implementation. `llms.txt` was chosen (at least initially) as a simple file in root to avoid requiring changes to HTML layout or HTTP server configuration. The `llms.txt` proponents argue a standalone text file is a low-friction solution: any site hosting static content can drop in a Markdown file with no risk of breaking site presentation.

Importantly, web search giant Google has weighed in on this proliferating ecosystem. In an July 2025 Search Engine Land report, Google's Gary Illyes (of the Search Central team) explicitly said **Google will not process llms.txt files**: "Google's AI Overviews rely on standard SEO; you don't need llms.txt or any special file" (Source: searchengineland.com). Illyes reaffirmed in public discussion that Google "doesn't support LLMs.txt and isn't planning to" (Source: searchengineland.com). Instead, Google instructs webmasters to *just use normal SEO* to be visible in AI-driven "AI Overview" features. In contrast, some smaller startup AI products (like OpenAI's engines or Claude) appear to be exploring or even actively reading these files. For example, one web developer reported that OpenAI's crawler was hitting his sites' /llms.txt endpoints every few minutes (Source: searchengineland.com). Thus at present, it seems llms.txt may be relevant for specialized AI tools, but not for mainstream search indexing.

In summary, **llms.txt occupies a new space**: it is explicitly intended not for search engines, but for AI agents. It complements rather than replaces `robots.txt` or `sitemap.xml`. It is *inspired by* these older conventions (hence sometimes called the "robots.txt for AI" (Source: www.released.so), but its guidance is of a different nature. Whether LLMs and companies will adopt this convention is a central question (addressed later), but technically it fills a unique niche: making complex site content easily consumable by generative AI.

## The Rationale: Why `/llms.txt` Might Matter

Understanding the importance of `llms.txt` requires examining the motivations and anticipated benefits from multiple angles: for content owners, for AI developers, and for end users.

**1. Control over AI interpretation:** The most often-cited benefit is giving website owners some control over how AI uses their content. In the current landscape, large AI models typically train on massive, uncategorized web scrapes (e.g. Common Crawl) or fetch pages ad-hoc without guidance (Source: privacyinternational.org). Authors and businesses have expressed concern that this process may misrepresent or misinterpret their content — or that AI may answer user questions without giving due "citation" or context. By providing `llms.txt`, a site can *highlight* the exact pages and data it *wants* AIs to read. This can ensure, for example, that product descriptions or legal terms are included, while unimportant pages (like navigate menus, login, or error pages) are left out. According to the proposal authors, this transparency can be a form of content rights management: websites can effectively signal which content they allow an LLM to "ingest" for answering queries (Source: llmscentral.com) (Source: www.released.so). In this view, llms.txt becomes a counterpart to the ongoing debate about AI training data and copyright. As Search Engine Land notes, content creators see it as "some assurance of increased control by the owner, in terms of what, and how much should be accessed" (Source: searchengineland.com).

**2. Improved AI answer quality:** When an LLM has direct access to a concise knowledge base, its generation quality improves. If an AI assistant is answering questions about your site or domain, you want it to have authoritative sources to draw from. Parsing raw HTML can yield context-free "hallucinations" or omissions. By contrast, a well-crafted `llms.txt` file summarizes key facts and links up to date information. Practitioners have reported that, after feeding an LLM the content listed in `llms.txt`, the AI provides more accurate and relevant answers about the site. For example, one practitioner tested an `llms.txt` file for a company called Enhance Media using three models (ChatGPT, Gemini, Claude) and found that *all three were able to correctly summarize the business from that file alone* (Source: www.linkedin.com). The file's structured format helped the models quickly home in on the salient points. Similarly, FastHTML's creators found that carefully curated context (via an expanded `llms.txt` file) produced "dramatically better results" from Claude and other tools than untargeted scraping (Source: news.ycombinator.com).

**3. Technical efficiency:** Large-scale crawlers (especially for smaller AI models) are resource-intensive. LLM companies must balance how often to re-scrape sites for fresh data. A `llms.txt` offering can serve as a **freshness beacon**: it may allow an AI crawler to check a single file for updates rather than crawling the whole site. Indeed, as reported in [33], at least one OpenAI system was polling developers' `llms.txt` every 15 minutes for updates . This kind of streamlined workflow can reduce unnecessary load on both the AI and the web servers. It can also ensure that the version of content the AI is exposed to is the official, flattened version provided by the site—not a partial or outdated scrap. In effect, `llms.txt` could serve as an "API" of sorts for static site content, albeit without the formal structure of an API call.

**4. Leveling the playing field:** Smaller sites and new startups may see llms.txt as a way to compete for attention in AI-driven search. Some analysts have drawn a parallel to early SEO strategies: in the web's infancy, small businesses used `robots.txt`, Meta tags, and sitemaps to stand out to search engines. Now, if AI agents become new "curators" of content, any site can use `llms.txt` to stand out to them too. This democratizing angle is explicitly mentioned in promotional materials: by adding `llms.txt` and even sharing it on platforms like GitHub, *"you'll shape how AI treats your content"* (Source: llmsly.com). The idea is that forward-thinking websites may gain a reputational advantage by being the first to partner with AI.

**5. Precedent of AI "robots":** Already, some AI tools present themselves as agents that crawl the web. For instance, Claude Projects (an IDE integration) can take documentation files into context. Such tools often require users to point them at key docs or data. `llms.txt` can automate that process. By offering a well-known anchor file, site owners can automatically enroll in these emerging AI ecosystems. It is similar to the early role of `robots.txt`: at first, few sites used it, but as Googlebot and others learned to check it, it became standard. Early adopters of `robots.txt` (circa 1994-95) did so to guide the AltaVista or Google crawlers. Today, designers of `llms.txt` hope the "architects of AI" (some leading AI teams) will do similarly. Indeed, the creators often highlight that developers from Anthropic are promoting `llms.txt` on their docs, and that companies like Mintlify built support for it (Source: www.released.so). In sum, `llms.txt` matters to its advocates because **it directly addresses a technical bottleneck of today's AI systems**. It promises a straightforward way to make the Web more "LLM-compatible," potentially making AI's job easier and answerable.

# Adoption, Industry Response, and Case Studies

How widely is `llms.txt` being used in practice, and who is paying attention? Since the idea first surfaced in late 2024, adoption has been limited and uneven, but certain clusters of activity are noteworthy.

First, **tech companies and documentation platforms** have shown interest. In November 2024, documentation platform Mintlify announced built-in support for `llms.txt` for projects published on their site (Source: www.released.so). This meant, practically overnight, *thousands* of software projects' documentation became `llms.txt`-accessible. The blog post by Jens Schumacher notes: "In one move, they made thousands of dev tools' docs LLM-friendly, like Anthropic and Cursor" (Source: www.released.so). Developer tool projects whose docs run on Mintlify (for example, many open-source libraries) thus acquired llms.txt files without individual action by maintainers. Similarly, some tech companies are explicitly creating `llms.txt`. In [15], Radu Stoian claims that **Anthropic** (the company behind the Claude AI) and unspecified others publicly requested llms.txt files for their sites: "AI leaders like Anthropic… have initiated it…they have built their models with the expectation of finding this file" (Source: www.linkedin.com). We have independently verified that https://www.anthropic.com/llms.txt (or the equivalent statically generated link) indeed exists and lists dozens of pages on Anthropic's site (Source: llmstxtgen.com).

Beyond developers, **consultancies and agencies** have begun recommending llms.txt. For example, a business-oriented blog author calls it "your new secret weapon" for AI optimization (Source: llmsly.com). Other SEO-focused websites and LinkedIn articles hail `llms.txt` as "essential for brands" in the AI era (Source: www.linkedin.com), giving it high-level visibility within marketing circles. A significant number of smaller companies and service providers (from SEO agencies to AI vendors) have blogged how to implement `llms.txt` on client sites. This enthusiasm is partly exploratory — many see AI content as the next frontier of visibility, and are treating llms.txt like a best practice to test.

However, when we scrutinize *actual usage*, the picture is mixed. A crowd-sourced directory of llms.txt files, [llmstxt.site], attests to hundreds of websites where `llms.txt` has been detected (Source: llmstxt.site). This directory lists dozens of example sites and their token counts. For instance, popular design tool **Framer** has an `llms.txt` with about 1,821 tokens (text size) (Source: llmstxt.site). The fintech company **Klarna** (in its documentation subdomain) has 17,387 tokens in its `llms.txt` (Source: llmstxt.site). Even a seemingly large content site, Weather.com (The Weather Company), is listed as having a (blank?) llms.txt (0 tokens) (Source: llmstxt.site), suggesting it might have created the file but left it empty. On the smaller scale, many personal, educational, and tech blogs have implemented `llms.txt`, occasionally with thousands or even hundreds of thousands of tokens. For example, an astrology blog "LookUpTheStars" reports a `llms.txt` with ~385,221 tokens (Source: llmstxt.site). At the other end, some `llms.txt` files are just a few hundred words (e.g. Ideanote.io had 1,106 tokens) (Source: llmstxt.site). Our survey of the llmstxt.site directory reveals widespread *experimental* adoption: companies of various sizes, from software products to niche ecommerce, have created these files (often converting existing sitemaps or manual link lists). Many appear to follow the spec format precisely, whereas a few have incomplete or ascending implementations (examples of parser tips are available in community forums).

To get a broader sense of adoption, two analyses have been reported by third parties. One is an "Industry Report" by a site called *LLMS Central*, which claims to have analyzed 2,147 websites across 15 industries in early 2025 (Source: llmscentral.com). Their headline statistics are that **68% of sites "allow" AI training** (with either fully open or selective policies), 23% "allow all," 45% have "selective policies," 18% block all, and only 14% have no `llms.txt` at all (Source: llmscentral.com). They interpret this to mean a majority of sites are publishing some guidance for LLMs. Notably, in their sample of tech & software companies (n=387), they report 95% have an explicit `llms.txt` policy of some kind (Source: llmscentral.com). These numbers, however, should be taken with caution. The report does not disclose how sites were chosen or whether they simply scraped for any mention of `llms.txt`. It is possible that their dataset is enriched for companies already engaged in AI/tech, which skews the percentages upward.

In sharp contrast, an SEO analytics firm *Rankability* published a monthly "LLMS.txt Adoption Report" focused on the **top 1,000 commercial websites** by traffic (Source: www.rankability.com). They found virtually **zero adoption**: 0.3% adoption rate (effectively 3 out of 1000) (Source: www.rankability.com). They state bluntly "Zero current adoption" (Source: www.rankability.com), with an extensive automated scan yielding almost no positive hits. By industry, their data shows 0.00% adoption across e-commerce, social media, finance, healthcare, government sectors, with only 0.73% adoption in the education sector (suggesting maybe 7 out of 1000 are universities or similar outliers) (Source: www.rankability.com). In short: among the world's largest sites, practically none implement `llms.txt` as of mid-2025. This implies the standard remains niche.

Why such discrepancy? It appears that adoption has been concentrated among **smaller or tech-oriented sites**, and virtually none among major mainstream brands. The top500-1000 list comprises global giants (Amazon, YouTube, etc.) with entrenched SEO teams; evidently, it has not yet penetrated those circles. By comparison, small-to-mid sites, knowledge bases and developer tools have flocked to it. The Rankability data suggests exactly one or two *isolated cases* in 1000 were found (likely small sites that ranked just into 1000). Meanwhile, the LLMS Central report likely sampled companies at least partially engaged in AI discussions, hence its higher adoption figures. This gap between "enthusiast community" and "mass market" will be important in assessing how much real-world impact `llms.txt` can have.

Given these figures, it is fair to say **`llms.txt` has spark but not (yet) flame**. It matters in certain ecosystems (especially software docs and SEO-agency commentary) but not broadly across the web. That said, adoption trends could accelerate if major platforms like Google or Microsoft's Bing decide to leverage it. Alternatively, it may remain an optional optimization for a subset of site owners. Next, we explore some detailed examples of `llms.txt` use, as well as reactions from AI tool developers.

## Case Study: Technical Documentation

One early and logical use case is software technical documentation. Developer docs often already generate HTML content from markup (e.g. Markdown) and generally strive to be machine and human readable. They also benefit greatly from precise answers. The FastHTML library discussed earlier is one example: its developers created `llms.txt` entries to assist developer-oriented AIs. Another prominent example is **Klarna's developer docs** (the European payments company). According to the llmstxt directory, Klarna's docs (hosted at docs.klarna.com) include an `llms.txt` with roughly 17,387 tokens (Source: llmstxt.site).

Similarly, a GitHub project "pgai/llms.txt" indicates that the Postgres AI (Timescale) project added an `llms.txt` file to its repository, suggesting an implementation in a real database product (Source: github.com). Corporate APIs, open-source libraries, and cloud platforms (the directory lists entries for AWS, Azure docs, etc.) have also begun adopting the format. These uses make sense: technical users are likely to benefit from having clear AI-readable docs summaries.

## Case Study: Pracademic & Services Sites

Not all adoption is in high-tech. For example, the SEO directory lists **HoodChefs** (a kitchen rental service) with 44,494 tokens (Source: llmstxt.site), and an **auto dealership** website "AutoChampion24" in Germany with 6,750 tokens (Source: llmstxt.site). These entries show that even small businesses see potential. "GalaxxiaMarketing" (a Brazilian marketing firm) has 676 tokens (Source: llmstxt.site), apparently introducing its services via llms.txt. Religion and spiritual sites, personal blogs, and e-learning providers have also been spotted. The existence of a site like "lookupthestars.com" with 385k tokens (Source: llmstxt.site) is notable: it appears to be an astrology-themed site that fully embraced the standard.

It is challenging to verify the business motivation for each ad-hoc `llms.txt`, but many likely did it out of curiosity or to experiment with SEO. Community contributions to llmstxt directories suggest WordPress plugins have been made to auto-generate llms.txt, and developers on forums mention times when their AI tutoring bots first saw llms-txt support.

### Industry Endorsements

Some major players have at least **acknowledged** the concept. Cloudflare's blog (May 2025) discusses how their new AI Gateway services integrate with other AI tools, although it does not directly mention llms.txt (Source: www.cloudflare.net). More relevant is Anthropic: their documentation portal now includes a visible "LLMS.txt" file link, and they have "posted on X" about supporting it (Source: www.released.so). In short, **AI-oriented companies** are at least curious. In contrast, mainstream tech or media firms have been silent. We are not aware of any reports of llms.txt adoption by Google, Amazon (beyond ones in the public directory), or Facebook.

### Metrics and Analytics

Few data exist on the *effectiveness* of llms.txt. One rough metric comes from a LinkedIn author who poked at Google Search Console analytics. He claimed that Google had already indexed an llms.txt file from a test site (Source: www.linkedin.com), though Google says it won't *use* them. Another trace cited is server logs: one webmaster noticed *OpenAI's* crawlers pinging his sites' llms.txt files every 15 minutes for freshness (Source: searchengineland.com). This anecdote suggests that at least some advanced search/AI services are paying attention.

Other metrics might include changes in query answers or referral traffic. As of this writing, such data are mostly not public. In theory, one could track traffic from AI chat interfaces (via special UTM tags or "referrals" from APIs), but few site owners have such tracking in place. Some SEO articles suggest using custom APIs to monitor LLM-driven traffic, but concrete examples are scarce (the golevels.com guide discusses it conceptually). Early signs in search results may also indicate usage. A LinkedIn post by an SEO consultant showed Google Search results highlighting an `llms-full.txt` file in results, hinting at indexing (Source: distinctly.co), but it's not clear if that is official or a glitch.

### Adoption By Region or Sector

The data from Rankability shows that **education** is the only sector with any measurable (0.73%) presence at the top sites (Source: www.rankability.com). This might be due to universities or scholarly projects experimenting with the format. In contrast, sectors like e-commerce, social media, finance, healthcare, and government had 0% in the top1000 (Source: www.rankability.com). The LLMS Central report (though less authoritative) indicates **technology/software** companies are leader in adoption, with "95% having explicit policies" within that segment (Source: llmscentral.com). This matches intuition: technology publishers are the earliest testbeds of AI tech.

## Criticisms, Concerns, and Alternative Perspectives

For balance, we must address reasons why `/llms.txt` may *not* catch on or could be problematic. Several criticisms have emerged from developers, SEOs, and skeptics. We organize them here:

**A. Duplication of Effort and UX Concerns:** Critics observe that if a site is already well-structured and has "help" or "about" pages, adding llms.txt may be redundant. As one Hacker News discussion pointed out: *"This isn't good UX for machines. This is a patch for bad UX to help LLMs... Some websites have the same patch for humans in the form of a 'Help' or 'About' section"* (Source: news.ycombinator.com). In other words, ideally a well-designed site should already make core info accessible, and a reader (human or bot) should find it naturally. If the site's actual content were simpler or more textual (e.g. via a "reader view"), an AI might not need llms.txt. This critique is essentially saying: "Fix the website, don't paper over its flaws." It also warns that llms.txt is a kind of shortcut that could discourage improving the underlying site design (like cramming content into an SEO block quote rather than authoring a usable interface).

**B. Limited Scope (Training vs. Inference):** It is important to clarify that llms.txt mainly affects the *inference-time* use of websites by AI, not initial model training. Many content owners want to control how their content is used to train new models (a legal and ethical debate), but llms.txt as specified does not directly enforce or register training permissions. It simply helps an LLM *fetch* content to answer queries. As Search Engine Land argues, the key differences revolve around indexing vs. usage: *"Robots.txt is all about managing crawling while the copyright discussion is all about how the data is used"* (Source: searchengineland.com). Critics may say: if a company doesn't want its site in AI outputs at all, llms.txt doesn't stop anyone (it only guides). Conversely, if the company already licenses content explicitly (e.g. with Creative Commons), llms.txt adds little. Konstantinos Zoulas's 2023 GEO

article suggests Creative Commons licenses (CC0, CC-BY, etc.) could govern AI use more directly than robots or llms directives (Source: searchengineland.com). This view implies llms.txt solves only the symptom (data discovery) not the heart of the content-rights issue.

**C. Lack of Standardization and Enforcement:** Currently, `/llms.txt` is a voluntary proposal without any formal RFC or registry. As Jeremy Howard himself admitted on Hacker News, it has not been registered under the IANA .well-known URI registry (a step required for official standard status) (Source: news.ycombinator.com). Without a formal decision or industry-wide endorsement, there's no guarantee software will reliably look for it. Critics point out that even robots.txt is not strictly enforced — it's a convention — and Google has shown it can ignore "robots.txt" if needed for legal reasons. With llms.txt even more in flux, some argue it might fizzle if key players stay on the sidelines. (Google's posture of ignoring it may already have dampened enthusiasm.)

**D. Potential for Misuse or Gameability:** As with any SEO-like signal, one might worry about spam or "gaming" llms.txt. In principle, a malicious site could create an llms.txt that contains misleading or malicious links, or bury tracker or ad URLs. However, because llms.txt does not automatically inject content into AI's training data, this risk is limited. It is more a risk that an unscrupulous site might stuff their llms.txt with irrelevant links just to push users (via AI answers) to them. The current spec does not specify any validation or rate limiting. How would an AI tool know if an llms.txt is legitimate? This is an unresolved question. In practice, since the format is human-readable and presumably curated, blatant abuse would likely be spotted and discredited by the community before it proliferated.

**E. Performance Impact on Websites:** Another concern (mostly hypothetical) is whether crawling and serving these potentially large text files could burden web servers. As noted, some `llms.txt` files run into the hundreds of kilobytes or even megabytes, comparable to a small HTML page. If an AI system polls them frequently (like every 15 minutes as one log indicated (Source: searchengineland.com), this could impose nontrivial load. Site operators should be mindful—though this issue is parallel to the pre-existing concept of "sitemap.xml" polling. Servers could always cache and throttle; it's a technical detail but one that must be implemented by web admins if `llms.txt` gains traction.

**F. Confusion Over Names and Versioning:** There is some ambiguity in terminology: the original proposal uses "`llms.txt`", but many posts (and LinkedIn articles) write it as "LLMS.txt" (with uppercase or plural LLMs). The community has generally settled on "llms.txt" (lowercase file name). Also, different tools talk about `llms-full.txt` (which contains the full concatenated text of pages) vs `llms.txt` (which lists links). This can confuse newcomers. Standardization or naming might evolve, but as of now this confusion may deter casual adoption.

**G. Alternative Approaches (No New File):** Finally, the most fundamental critique: *Do we even need a new file at all?* Some SEO experts argue that the same goals could be met by revitalizing older ideas. For instance, openAI's early discussions mentioned using "noindex" or "nofollow" in robots to differentiate regular search vs AI use (Source: searchengineland.com). Others propose entirely in-band signals: e.g. Google (in mid-2023) suggested simply using normal links and SEO practices so that AI (like Google's own Overviews) naturally find content (Source: searchengineland.com). There is also the concept of an HTTP header or element that identifies a file or format for LLMs, rather than a raw text file. Some commenters say this would be more semantically "web-like" than inventing yet another file type. Pro-llms advocates generally respond that nothing prevents using multiple approaches (header *and* llms.txt), but this remains an area of discussion.

In sum, the criticisms focus on **practicality and necessity**: If Google (and Bing) get all content via old methods, llms.txt may be overkill. If AI developers could just scrape HTML better or use embeddings from existing search indexes, perhaps they don't strictly need it. At the same time, supporters point out these issues have not deterred initial experiments or standards formation. Whether these concerns prove fatal or surmountable will likely depend on concrete usage and community momentum.

# Data and Analysis

A thorough analysis of `/llms.txt` requires not just descriptions but data-driven insights. However, as of mid-2025 the ecosystem is still nascent. Below we summarize the available data and quantitative findings:

- **Adoption in Web traffic rankings:** The Rankability study is one of the few publicly reported analyses of adoption. It surveyed the top 1,000 most-visited websites (globally) as of mid-2025 and found **0%** usage of `llms.txt` (only ~0.3% by one count, rounding to 0%) (Source: www.rankability.com). Breaking that down by sector, it reported 0.00% adoption in every major industry category (e-commerce, social media, finance, etc.), except a tiny blip of 0.73% in Education (Source: www.rankability.com). This suggests that, among the Web's heavyweights, virtually none have implemented `llms.txt`. In

practical terms, if you Google any big site (e.g. Wikipedia, CNN, Amazon), you will find no `llms.txt` unless someone explicitly set one up just to test. (Notably, Rankability's definition of "adoption" likely required an HTTP 200 response for /llms.txt. Some sites returning 404 or error would count as non-adoption.)

- **Adoption among surveyed sites:** In contrast, a different analysis of a broader set of 2,147 websites (the "LLMS Central" report) claimed 86% of sites had **some llms.txt content** (68% allowing AI training fully or selectively, and only 14% having none) (Source: llmscentral.com). Their methodology isn't fully transparent, but they grouped site policies as "Allow all", "Selective", "Block all", or "No file". Seeing a category like "Allow all" (23%) implies these sites have an llms.txt explicitly stating to allow LAI usage. If taken at face value, this report suggests **overtwo-thirds of mid-sized sites** in their sample published an llms.txt. It also finds tech companies especially avid: 95% of technology/software companies they surveyed had an llms file (Source: llmscentral.com), vs smaller percentages in other industries. However, without knowing their sample selection, this may reflect a self-selection bias (maybe they scraped sites that already mentioned AI on their blogs).

- **File sizes and content:** Looking at actual llms.txt contents, we see tremendous variation. The example in Table 2 below shows some representative token counts for a few sites (from the llmstxt.site directory). These kinds of numbers give a sense of scale. Notably, some technical documentation sites result in huge llms files: for instance, *M-Source* (a database company) has **328,716 tokens** listed (Source: llmstxt.site), and *LookupTheStars* has **385,221 tokens** (Source: llmstxt.site). (For context, GPT-4's context limit is around 32k tokens, so a single llms.txt of 300k tokens would need to be chunked.) Others are token-lighter: Ideanote.io's llms.txt is 1,106 tokens (Source: llmstxt.site), HoodChefs 44,494 tokens (Source: llmstxt.site), Framer 1,821, Klarna 17,387, etc. An extreme outlier is *X-CMD*, whose llms-full file is 590,515 tokens (Source: llmstxt.site) (implying a colossal site or possibly a quirk of how it's generated). The variability indicates that sites interpret how much to include differently.

- **Crawling & Traffic Insights:** There is little public data on traffic. One table from the SEO reporting site [33] highlights that Googlebot requests for llms.txt happen *zero* times ("Google won't be crawling your LLMS.txt" (Source: searchengineland.com). By contrast, user Ray Martinez reported in his site logs that "OpenAI crawls my LLMs.txt file on a few sites… pinging our servers every 15 minutes looking for freshness" (Source: searchengineland.com). This log analysis suggests that, at least for his sites, **OpenAI's systems are actively checking llms.txt often** (perhaps assuming they should). Google's John Mueller similarly said in an earlier Search Console hangout that "no AI system is currently using the LLMS.txt file" (Source: searchengineland.com) (quote from seroundtable). In sum, the only empirical insight we have is anecdotal: Google search ignores it, some AI labs poll it.

- **SEO performance correlation:** No credible aggregated data exist linking llms.txt with improved search ranking or traffic. Google explicitly says normal SEO is adequate (Source: searchengineland.com), implying they found no advantage. It remains to be seen whether, for example, inclusion of llms.txt will positively affect snippets or "answers" in AI chat interfaces. In principle, if an AI assistant directly cites llms.txt content, a savvy marketer will try to detect that and optimize accordingly. But as of mid-2025, this remains hypothetical.

- **LLM Tool Support:** Beyond Google, notable LLM products have begun acknowledging llms.txt. *Anthropic* (Claude) documentation includes it; *LangChain*'s MCP (multi-context plugin) supports reading llms.txt from IDEs (Source: github.com). Some open-source LLM-based chatbot frameworks now have boilerplate to look for llms.txt. The very existence of a GitHub repo (AnswerDotAI/llms-txt) and automated CI tests indicates developer interest. On the other hand, major platforms like ChatGPT (OpenAI front-end) have not announced formal support, aside from backend indexing. Analyst reports from Distinctly (SEO news) have noted a screenshot of ChatGPT pulling content from an "llms-full.txt" (Source: distinctly.co), but details are lacking and this might be a one-off.

These data points paint the picture: **emerging but minor**. Dozens or hundreds of smaller sites have llms.txt, but no critical mass. If adoption were charted over time, we might see a slow rise among mid-tier sites in late 2024 through 2025, plateauing. A tipping point would likely require one or more dominant AI platforms to declare "yes, we use llms.txt." Otherwise, it may remain a niche best practice.

Below is a table summarizing some adoption statistics and examples:

| METRIC / SITE CATEGORY | VALUE / EXAMPLES | SOURCE |
|---|---|---|
| Top-1000 sites using llms.txt | ~0% (0.3%) | (Source: www.rankability.com) |
| Tech/Software companies (surveyed) | 95% (sites in those categories have llms policies in one report) | (Source: llmscentral.com) (Source: llmscentral.com) |
| Allow-all (all content open) | 23% of sites (per one report) | (Source: llmscentral.com) |
| Selective policies (some pages) | 45% of sites | (Source: llmscentral.com) |
| Block-all (no AI use allowed) | 18% of sites | (Source: llmscentral.com) |
| No `llms.txt` file | 14% of sites | (Source: llmscentral.com) |
| Example sites with llms.txt | Framer.com (1,821 tokens), Klarna docs (17,387), M-Source (328,716) (Source: llmstxt.site) (Source: llmstxt.site) | (Source: llmstxt.site) (Source: llmstxt.site) |
| Largest reported llms size | ~385,221 tokens (lookupthestars.com) | (Source: llmstxt.site) |
| OpenAI crawling frequency | ~every 15 minutes (site log) | (Source: searchengineland.com) |
| Googlebot llms.txt requests | None reported; Google says it won't crawl llms.txt | (Source: searchengineland.com) |

*Table: Selected figures relating to llms.txt adoption and usage.*

## Perspectives and Expert Opinions

To fully grasp the stakes of `/llms.txt`, we consider what various experts and stakeholders have said—sometimes loudly—about the proposal.

- **Jeremy Howard (Answer.AI, fast.ai):** *Proponent and proposer of llms.txt.* He argues primarily from the perspective of developer usability. In discussion threads, Howard emphasized that the aim is to help "end-users use the information on websites with the help of AI" (Source: news.ycombinator.com). He gave concrete examples: when he released the FastHTML library, many potential users complained AI tools (cursor, etc.) could not answer questions about it because the models post-date their knowledge. His solution: manually curate the documentation once in an llms.txt so AI tools have it readily at inference time. Howard frames llms.txt as an end-user/community aid rather than a scraping concern: "llms.txt isn't really designed to help with scraping; it's designed to help end-users use the information on web sites with the help of AI" (Source: news.ycombinator.com). He also stresses that providing llms.txt saves *everyone* effort: instead of every engineer individually picking context for prompts, the site owner does it once. In interviews and blog posts, he frequently mentions developer-doc use cases, and the fact that many fast.ai/nbdev docs now auto-generate markdown to satisfy this need (Source: github.com).
- **SEO/Marketing analysts (e.g. SearchEngineLand, Expecting SEO Agencies):** Broadly, SEO publications have taken a cautiously optimistic view. The March 2025 SEL article by Roger Montti surveyed llms.txt and noted both "interested content creators" and "detractors" (Source: searchengineland.com). Montti's stance is neutral-to-curious; he explains the spec and suggests it "increases control by owner". Roger highlights the resource saving angle (LLMs focus on intelligence, not crawling).

Meanwhile, others in the SEO community hype llms.txt as a must-have for brands. For example, Radu Stoian's LinkedIn piece bluntly titles it "non-negotiable for your brand" (Source: www.linkedin.com). Such pieces promise improved brand narrative and even claim Google is indexing llms.txt now. However, as an unvetted blog, these should be read with skepticism. More measured voices (outside SEL) suggest llms.txt is an incremental "AI SEO" technique : a possible optimization but unlikely to overtake traditional SEO (Source: llms-txt.io).

- **Google Search Engineers:** The clearest statements have come from Google itself, albeit indirectly. At a Google Search Central event in July 2025, Gary Illyes (Search Analyst) made it explicit: *"To get your content to appear in AI Overview, simply use normal SEO practices... It also said Google won't be crawling the LLMS.txt file."* (Source: searchengineland.com). In effect, Google's message is: *Ignore llms.txt in terms of search ranking – we don't use it.* This was echoed by John Mueller, who said in a Webmaster Hangout that "no AI system is currently using the LLMS.txt file" (Source: searchengineland.com). These assertions mean that, from Google's perspective, llms.txt has no bearing on SEO. It may discourage publishers who primarily care about Googleability. It also raises a bigger question: even if llms.txt is beneficial to your content's encounter with *some* AI, if that AI is not the one dominating searches (Google Search), the impact on real traffic might be small.

- **OpenAI (ChatGPT developers):** OpenAI has not publicly commented on llms.txt, but limited evidence suggests they have at least *tested* or allowed its use. The log analysis by Ray Martinez is unimpeachable evidence that some OpenAI infrastructure is polling llms.txt for changes (Source: searchengineland.com). This suggests OpenAI's agents have recognized llms.txt in the wild and treat it as a "freshness endpoint." However, OpenAI spokespersons have not announced any policy stance. Anecdotally, users of tools like ChatGPT's "Browse with Bing" plugin or third-party agents anecdotally try to leverage llms.txt, but no official documentation is available.

- **Anthropic (Claude developers):** Anthropic is widely believed to support llms.txt. Their docs team added it early, and Anthropic engineers have signaled interest in standardization. Claude Projects (Claude's code IDE plugin) treats llms.txt as a first-class citizen: users loading a codebase can specify an llms.txt. One community snippet on GitHub shows instructions for configuring Claude Desktop/Cursor to read llms.txt (Source: gist.github.com), implying built-in support. Kohl Marcus (in Distinctly news) mentioned that "Aimee Jurenka shows ChatGPT accessing content from an llms-full.txt file" (Source: distinctly.co), so presumably that was via Anthropic's frameworks. All this indicates at least advanced AI products (like Claude) are taking llms.txt seriously.

- **Academic and Privacy Experts:** Organizations concerned with data privacy note that llms.txt touches on the scraping narrative. Privacy International, in an explainer about LLMs, underscores that *"the more written language [LLMs] can get hold of, the better"* and that web scraping is often "indiscriminate" (Source: privacyinternational.org). While they don't mention llms.txt specifically, the implication is that anything which makes scraping more targeted (i.e. guided by owners) could align with data governance. No formal privacy law recognizes llms.txt, but advocates like Jay Graber (Bluesky CEO) who lead AI creator rights debates have pointed out that llms.txt and other initiatives (like the "Bletchley Declaration") are part of emerging norms for data control in AI. In short, some see llms.txt as a constructive gesture toward respecting content ownership, even if it's non-binding.

- **Critics and Pragmatists:** Many coders and SEOs approached llms.txt pragmatically. On Hacker News and blogs, commenters voiced skepticism: one noted that if a site's UX is good, an "instructions page" could suffice and llms.txt would be unnecessary (Source: news.ycombinator.com). Others said that maintaining an extra file is overhead; they'd rather rely on `rel=search` or API-based approaches. From a standards perspective, one commenter pointed out that perhaps a `<link rel="llm">` tag or HTTP content-type negotiation might be more elegant than a text file (Source: news.ycombinator.com). These suggestions reflect a desire for solutions that integrate smoothly into existing Web architecture, rather than adding a parallel silo.

Despite these mixed views, **the common thread** is: `llms.txt` forces the question "Should the Web adapt for AI?". Many interviewed voices pride themselves on being early adopters. Fans argue it lets websites *join the conversation* rather than be passive data mines (Source: www.linkedin.com), while detractors say it disrupts the Web's uniform interface. Ultimately, most see it as an *experiment*: an idea worth testing now, with community feedback guiding whether it becomes a de facto standard or fades.

## Implementation Considerations and Tools

For a website owner considering adding `llms.txt`, practical questions arise: *How to create it? What content to include? How to maintain it?* Fortunately, several tools and guides have emerged to address these.

- **Guides and Examples:** The llmstxt community site (llmstxt.org) features example llms.txt files and a step-by-step guide. There are also numerous blog articles and GitHub repositories with sample llms.txt implementations. Key advice includes: start with the homepage/title, write a succinct summary (about 1–3 sentences) in a blockquote, then list crucial pages. Some SEO blogs recommend adding company info (contact, address), FAQs, developer docs, product pages – basically everything a helpful AI might need to answer user queries (Source: llmsly.com) (Source: golevels.com). It's often suggested to keep the file under a few megabytes; one post mentioned that llms.txt files can range from a few KB to hundreds of KB (Source: searchengineland.com). The format is flexible: you can use images (as links), bullet points, or short paragraphs. Some sites even break llms content into multiple files: the *llms-full.txt* variant can contain whole sections of text if needed.

- **Existing Tools:** Several open-source tools help generate or validate llms.txt. For example:

  - **llms.txt Generator (llmstxtgen.com):** A web app where you paste your sitemap or URL list; it crawls and outputs a draft llms.txt in seconds. The screenshot [10] shows one tool's auto-generated output (for anthropic.com).
  - **CLI Utilities:** The GitHub repo (AnswerDotAI/llms-txt) includes scripts like `llms_txt2ctx` which can combine llms.txt and linked markdown into a machine-consumable context file (Source: github.com). Others (like Firecrawl's tool referenced in [66]) can crawl and assemble content into markup lists.
  - **CMS Plugins:** There are plugins for WordPress and other CMS that generate llms.txt from site menus or posts (as hinted by [59]). These allow dynamic updates as site content changes.
  - **IDE/LLM Integrations:** Tools like LangChain's `mcpdoc` can pull an llms.txt automatically when setting up AI, so developers don't have to fetch it manually (Source: github.com). This shows llm frameworks starting to recognize the file.

- **Maintenance:** Given sites change, llms.txt needs updates. Unlike sitemap.xml (which can be automated), llms.txt is more manually curated. However, some workflows create it from existing site data: e.g., a script can scan navigation menus to list URLs, or compile README files. The Ethereum docs project, for instance, uses a CI process to rebuild llms.md whenever docs change (as part of its static site generation). Broadly, it is recommended to review llms.txt whenever major site content changes, since stale links or summaries could mislead AI. Monitoring involves just checking uptime of that single file (e.g., site health checks).

- **Hosting and Performance:** As with any static asset, best practice is to serve llms.txt with caching enabled (HTTP cache-control) and gzip compression, since it is typically text. Large llms.txt files (hundreds of KB) can weight down bandwidth if crawled too frequently, so proper caching helps. Some have suggested hosting llms.txt on a CDN or exposing it via .well-known/llms.txt so proxies can cache it globally.

## Case Studies in Depth

**FastHTML (Hypermedia Framework):** The FastHTML project's experience is illustrative. FastHTML is a small library for creating APIs and docs. Its developers recognized that typical language models (like Claude) had no knowledge of FastHTML (it was released after their training cutoff). To compensate, they authored an llms.txt for their documentation site. Then, using `llms_txt2ctx`, they generated two versions of context files: `llms-ctx.txt` (core content) and `llms-ctx-full.txt` (extended with optional links) (Source: github.com). This allowed them to feed Claude a concise but complete view of the docs whenever answering questions. The outcome: they reported dramatically better AI-assisted answers in their IDE and documentation bots, without each user having to manually copy links. This demonstrates `llms.txt` serving the "long tail" of content (FastHTML's docs were not indexed by Google, as per [4]). Their case shows how a modest project can leverage llms.txt to make itself "AI searchable" from day one.

**Anthropic (AI Company):** Anthropic's adoption of llms.txt is more symbolic than case-specific. As a major AI company, they arguably have less need to be AI-findable, but they have nevertheless created llms.txt for transparency and community signaling. Their llms.txt lists introductions to their products (Claude), research papers, developer channels, and more (the output [10] shows pages like "Claude in Slack", "API", "Customers"). Their participation lends credibility: an industry leader including llms.txt suggests it's worth taking seriously. It also likely feeds back into Anthropic's own models (if they index it internally).

**Academic Institution (example):** Some universities have large websites with course catalogs, research, etc. One example is "Juris Education" which has a sizeable llms.txt listed (22,885 tokens) (Source: llmstxt.site). The rationale may be to help prospective students or AI tutors / chatbots collate course info quickly. Many universities experimented with AI portals for student Q&A, and llms.txt could serve as a backend resource.

**Government and Regulations:** As yet, there seem to be no official government guidelines on llms.txt. However, it resonates with policy debates. For example, the EU's Copyright Directive article on text-and-data mining provides exceptions for research, implying websites wouldn't need to explicitly opt-in for that use if it's in scope. LLms.txt sits in a grey area: it is voluntary metadata for AI data use, not a binding license. Some policymakers advocate more enforceable mechanisms (e.g. web scraping bot laws). No known government has mandated anything like llms.txt.

# Implications and Future Directions

Looking forward, the success or failure of `llms.txt` will likely hinge on a few key factors:

- **AI Platform Adoption:** If major AI models or tools come to recognize and trust llms.txt, its adoption could spike. For instance, if OpenAI officially supported it (e.g. via ChatGPT instructing GPT on an llms.txt link), or if Google changed course and indexed llms.txt, that would create a strong incentive. Conversely, if AI developers prefer to rely on search indexes or embeddings (like how Bing Chat already uses search results under the hood), the demand for llms.txt may stay limited. The fact that Google currently dismisses it suggests that mainstream "AI search" will be slow to embrace it. But the landscape can change rapidly: last we checked (June 2025), Google said normal SEO was enough (Source: searchengineland.com), but a year later that could flip if user behavior shifts towards AI summaries.
- **Tool and Framework Ecosystem:** Growth of developer tools around llms.txt could make it easier to adopt. For example, if GitHub Pages automatically generates llms.txt, or if Wordpress and other CMS include it by default, a flurry of new sites might be "llms.txt-ready" overnight. We've already seen the beginnings: a WordPress plugin exists, some static site generators have add-ons. If major content management systems integrate support, adoption could climb regardless of the big search players.
- **Standardization:** Moving from proposal to standard normally requires consensus and registry. The authors hinted at possibly registering it as a well-known URI (e.g., `/.well-known/llms.txt`) if the standard takes hold (Source: news.ycombinator.com). Such a move would make orientation easier for bots. Additionally, publishing an RFC or W3C note could cement the format. If llms.txt gets formal backing, that could signal "official status," encouraging wider buy-in (much as RSS became ubiquitous once standardized).
- **Alternate Approaches:** It's possible that better solutions emerge. For instance, Google might develop its own "AI sitemap" or meta tags to control AI indexing, rendering llms.txt obsolete. Or AI assistants could use contextual signals (schema.org markup, Knowledge Graph data, voice assistant schemas) to glean information more semantically. There is an ongoing discussion about standards like SERP features or the "AI prompt hints" embedded in HTML. In the worst-case scenario, llms.txt could become *one* of many similar proposals, and perhaps get superseded by a more elegant protocol.
- **Regulatory Influence:** If regulators require AI companies to respect robots.txt (as part of scrapers regulation), a logical extension might be to require respecting llms.txt directives. This could happen through industry self-regulation or law, especially as debates about AI training data and copyright intensify. For example, if the EU or a country legislated that AI systems must honor website owners' published content use preferences, they might explicitly mention llms.txt as a recognized channel. This is speculative but within the realm of emerging AI governance.
- **Networks Effects on Content Discovery:** We are only at the early stages of "AI-driven content discovery." If one or two popular AI assistants start defaulting to llms.txt listings, users might start seeing it indirectly. For instance, if Gemini or Claude's answers regularly cite content from an llms.txt page, savvy content teams will notice and optimize their files. This is similar to how SEOs reacted when featured snippets started pulling from particular HTML structures (they then modified their content to feed snippets). Over time, good llms.txt practice could yield partial AI-SEO benefits not captured in traditional metrics.
- **Community Best Practices:** The `llms.txt` ecosystem itself will evolve through shared experience. As early adopters publish their experiences, community best practices will develop. The GitHub and blog resources are already documenting Do's and Don'ts (for instance, suggestions on how to structure blockquotes so they don't confuse an LLM). Over months, we expect linting tools for llms.txt to appear (checking for broken links, clarity, etc.). There may also emerge versioning conventions (like how robots.txt has no official version, llms.txt could either fix the spec or allow variations like llms-full.txt).

In conclusion, **the future of `llms.txt` is open-ended**. Many observers have noted that no single piece of technology can guarantee how AI evolves — whether the "content behavior" sector consolidates around just publishers (like llms.txt) or remains decentralized. For now, llms.txt sits in a niche but active corner of the AI web. If it catches on, it could lead to a new layer of web-

file standards; if not, it may quietly recede as an interesting experiment.

## Conclusion

Our investigation of `/llms.txt` finds that it is a well-defined proposal with specific aims: to make websites more accessible to large language models by way of a human-created map of content. The technical specifications (using Markdown, lists of links, etc.) are clear and relatively easy to implement. Early case studies in software documentation have shown that llms.txt *can* improve AI agents' performance on niche tasks (Source: searchengineland.com) (Source: www.released.so). Yet at the same time, there is an equal measure of skepticism. Major search engines have so far publicly proclaimed they will ignore this file (Source: searchengineland.com), and empirical scanning suggests mainstream sites have not yet adopted it appreciably (Source: www.rankability.com).

Does it matter? For now, the answer is: *It depends on your priorities*. If you are a technology publisher, developer, or SEO-savvy marketer who wants to experiment with every edge optimization in the AI era, `llms.txt` seems worth exploring. It imposes relatively little cost, is reversible, and if AI tools begin to support it extensively, you'll have gotten ahead of the curve. It particularly matters for domains where AI-powered Q&A may drive technical support or user onboarding: developer docs, APIs, product manuals, etc.

However, if you're focused solely on traditional search or you have limited resources, then `llms.txt` may be considered optional. The consensus from Google's SEO team is that "normal SEO" covers being found in AI results (Source: searchengineland.com). Organizations disinterested in AI training of their data (or opposed) might prefer more concrete legal mechanisms (licenses, robots blocks) instead of a friendly list. As the LLMS Central report implied, many content owners see `llms.txt` as part of **AI training transparency** – but whether an AI actually respects it (or compensates for it) remains largely untested.

Looking ahead, the most immediate effect of llms.txt is to spark vital conversations among webmasters about content design for AI. By trying this new tool, the community can discover where LLMs succeed or falter when digesting real sites. It informs both sides (site and AI developers) about what works. In that sense, llms.txt has already had some impact: it made AI trainers aware of context-window issues, and made SEO experts aware that search engines are not yet AI agents, etc.

Ultimately, the narrative around llms.txt echoes broader discussions on the future of the Web: Will content creators exert explicit control over AI use of their data, or will the Web remain a passive text corpus? Will we see an "AI web" with new mini-standards layered on HTML (much as there are now AJAX and JSON conventions), or will AI simply layer on top of existing infrastructure (semantic annotations, improved crawling)? The jury is still out.

What is clear is that **`llms.txt` matters to the extent that the industry and community decide it does**. If one sees it as analogous to how `robots.txt` and `sitemap.xml` gained traction, then its importance will grow as soon as enough content and enough AI systems converge on it. It is still early days, and for every under-the-hood technical benefit claimed, there are equally weighty concerns about necessity and viability.

In our view, `llms.txt` is a proactive and constructive experiment: it tries to preempt AI-related miscommunication on the web. Our research suggests that it is a well-intentioned solution that addresses real technical challenges (Source: searchengineland.com) (Source: www.released.so). Its future success will depend on both technical uptake (by AI platforms) and community adoption (by site owners). We endorse its continued exploration – after all, a negligible downside approach combined with even a small upside in AI fidelity seems a worthwhile gamble. Whether it becomes part of the standard toolkit for the internet, or just a footnote in the history of Web evolution, only time (and data) will tell.

**References:** All claims and figures above are drawn from the sources cited in the text (Source: searchengineland.com) (Source: searchengineland.com) (Source: llmscentral.com) (Source: www.rankability.com) (Source: llms-txt.io) (Source: www.released.so) (Source: llmstxt.site) (Source: www.kdjingpai.com), and from additional industry reports and expert commentary as detailed. Each citation identifies the source of the information described.