Annotated llms.txt template

llms.txt is a Markdown file published at /llms.txt that introduces a website and lists the pages most useful to a language model. The proposal comes from Jeremy Howard (Answer.AI) and dates from 3 September 2024; version 2 of the specification is from 10 August 2026.

The only mandatory element is the H1 title. Everything else is optional and recommended, and that is the part which decides whether the file is any use.

BEFORE YOU PUBLISH IT

Google says it does not support it. In July 2025, according to Search Engine Roundtable, Gary Illyes said Google does not support llms.txt and has no plans to; the official 2026 guide on optimising for generative AI features rules out any need for machine readable files, AI text files, markup or Markdown. In the sources I have read, the other major AI companies do not confirm that the file affects citations or ranking. It mainly helps the agents and tools that read it while working on a site: publish it for that, not for visibility on Google.

The full template

This is the sample file published in the aioverview.it guide, for an Italian company with made-up data. Copy it and replace the content: the shape holds, and I explain the lines one by one after it.

```
# Esempio Serramenti Srl

> Produttore di serramenti in PVC e alluminio con sede a Brescia.
> Vendiamo a privati e imprese in Lombardia e Veneto, con posa in opera
> e pratiche per le detrazioni fiscali.

Le schede prodotto riportano trasmittanza termica, classi di tenuta
e prezzi indicativi aggiornati. Le guide sono firmate dal nostro
ufficio tecnico e riportano la data di revisione.

## Prodotti

- [Finestre in PVC](https://www.example.it/finestre-pvc.md): gamma, valori di trasmittanza,
colori disponibili
- [Finestre in alluminio](https://www.example.it/finestre-alluminio.md): gamma a taglio termico
e confronto con il PVC
- [Listino indicativo](https://www.example.it/prezzi.md): prezzi medi al metro quadro, posa
inclusa

## Guide

- [Come scegliere una finestra](https://www.example.it/guide/scegliere-finestra.md): criteri
tecnici e domande da fare al fornitore
- [Detrazioni per i serramenti](https://www.example.it/guide/detrazioni.md): requisiti e docu-
menti necessari
```

```
## Azienda

- [Chi siamo](https://www.example.it/chi-siamo.md): storia, certificazioni, sede e area servita
- [Contatti](https://www.example.it/contatti.md): orari, telefono, richiesta di sopralluogo
```

Line by line

On the left the piece of the file, on the right why it is there and what changes if you take it out.

```
# Esempio Serramenti Srl
```

WHY IT IS THERE

The H1 with the name of the site or project. It is the only element the specification calls mandatory, and it settles the question "whose file is this".

IF YOU REMOVE IT

The file no longer follows the specification. Whoever reads it finds a list of links with no idea who they belong to.

```
> Produttore di serramenti in PVC e alluminio con sede a Brescia.
> Vendiamo a privati e imprese in Lombardia e Veneto…
```

WHY IT IS THERE

The summary blockquote: what the company does, who it sells to, where. Two or three lines. It is optional, but it is the only place in the file where you state the scope.

IF YOU REMOVE IT

The links remain. A reader has to infer the sector and the area served from page titles, and infers them badly: a firm in Brescia selling across two regions becomes a generic firm.

```
Le schede prodotto riportano trasmittanza termica,
classi di tenuta e prezzi indicativi aggiornati…
```

WHY IT IS THERE

The free paragraph after the blockquote says how the content is made: what the pages report, who signs them, whether they carry a revision date. It is for whoever has to decide whether to trust them.

IF YOU REMOVE IT

You lose the one line that talks about the quality of the pages rather than their topic. The file stays valid.

```
## Prodotti
```

WHY IT IS THERE

H2 sections group the links by kind: products, guides, company. The specification has them for exactly this.

IF YOU REMOVE IT

The links flatten into a single list. With ten entries you survive, with thirty you do not: someone after the price list has to scroll the lot.

```
- [Finestre in PVC](https://www.example.it/finestre-pvc.md):
gamma, valori di trasmittanza, colori disponibili
```

WHY IT IS THERE

The entry itself: an absolute link and, after the colon, the note saying what is on the page. The note is the part that makes the choice.

IF YOU REMOVE IT

Without the note you are left with a title. "PVC windows" does not say whether the page holds the technical values or only photographs, so the reader opens it at random.

```
https://www.example.it/finestre-pvc.md
```

WHY IT IS THERE

The .md extension points to the Markdown copy of the page. The specification recommends offering one, because a model reads it without menus, banners and scripts.

IF YOU REMOVE IT

You point at the HTML and it still works. The .md copies pay off only if the CMS generates them on its own: keeping them by hand costs more than it returns.

```
## Azienda

- [Chi siamo](…): storia, certificazioni, sede e area servita
- [Contatti](…): orari, telefono, richiesta di sopralluogo
```

WHY IT IS THERE

The section that describes the company rather than the catalogue. Address, certifications, opening hours and contact details are what an agent looks for to answer a practical question.

IF YOU REMOVE IT

The file only describes what you sell. Questions about who you are, where you are and how to reach you go unanswered, and somebody else answers them.

```
User-agent: GPTBot
Disallow: /riservato/
```

WHY IT IS THERE

This line does not belong in the file. Permissions live in robots.txt, bot by bot: llms.txt holds neither permissions nor prohibitions.

IF YOU REMOVE IT

Nothing: writing it has no effect at all. A line saying "do not use this content for training" inside llms.txt is decoration.

The .md versions

If you publish the .md versions, declare them in the HTML of the matching page. That way whoever looks for them finds them without guessing the address.

```
<link rel="alternate" type="text/markdown"
href="https://www.example.it/finestre-pvc.md">
```

Remove this line and the .md copies exist but nobody discovers them, apart from whoever reads llms.txt.

The mistakes I see most often

- Writing prohibitions into it. Bots that follow rules read robots.txt; for OpenAI, Anthropic and Perplexity the instructions go through user agents such as GPTBot, ClaudeBot or PerplexityBot.
- Using it instead of the sitemap. Google discovers URLs through internal links and sitemap.xml, and llms.txt plays no part in that.
- Describing a company the pages do not confirm. It stays a statement by the owner, worth about as much as the meta keywords tag.
- Buying it as a service with a promised percentage of extra citations. In the sources I have read, the engines do not confirm that effect.
- Letting it age. An index with broken links helps nobody: update it when the pages change.

Where to start

Where to start: pick 10 to 30 pages that genuinely describe the company, write the file, publish it at the root so that /llms.txt returns 200 as plain text, and keep it maintained. Do not measure its effect on citations: look in the logs to see who requests it, and keep measuring visibility with Search Console, Bing Webmaster Tools and a fixed set of questions.

Davide Filiaggi, aioverview.it · published by RankWit AI · the current version lives at aioverview.it/en/downloads/modello-llms-txt

Print it, pass it to colleagues, use it in a classroom. If you quote it, quote the source.