

OVERALL SCORE

32 /100 Poor

Executive summary

This site is difficult for AI tools to read right now. Key strengths include sitemap and structured data, while homepage access and crawler policy need attention. Recommended next step: remove AI crawler Disallow: / rules or replace them with narrower path-level restrictions for private content only.

Recommended next step

- 1. Remove AI crawler Disallow: / rules or replace them with narrower path-level restrictions for private content only.
- 2. Remove AI crawler Disallow: / rules or add narrower Allow/Disallow rules if AI crawlers should be able to discover public content.
- 3. Publish /llms.txt as text or markdown with more than 200 characters, markdown headings, and at least one absolute URL.

Signal breakdown

Crawlability Fail 0/20

Remove AI crawler Disallow: / rules or replace them with narrower path-level restrictions for private content only.

Robots.txt Fail 0/15

Remove AI crawler Disallow: / rules or add narrower Allow/Disallow rules if AI crawlers should be able to discover public content.

llms.txt Fail 0/15

Publish /llms.txt as text or markdown with more than 200 characters, markdown headings, and at least one absolute URL.

Sitemap	Pass	10/10
The sitemap.xml file is valid and contains URL entries.		
Markdown support	Fail	0/15
Add content negotiation for Accept: text/markdown on the homepage and return a markdown representation with Content-Type: text/markdown. Keep the HTML response for regular browser requests.		
Semantic HTML	Warn	7.1/10
Avoid skipped heading levels so sections progress from h1 to h2 to h3 without gaps. Add missing semantic elements: main, article.		
Structured data	Pass	10/10
Valid JSON-LD structured data was found with core Organization or WebSite schema.org types.		
Content signals	Pass	5/5
Consider adding Content-Signal HTTP header, AI-specific head meta tags, robots noai/noimageai directive so AI systems can consistently discover content usage preferences across robots.txt, HTTP headers, and HTML metadata.		

Suggested fixes

llms.txt

MARKDOWN

Newsletter Surf - Discover great newsletters

> A curated directory of newsletters worth your inbox. Browse hand-reviewed newsletters by category - tech, business, culture, and more.

This llms.txt file summarizes the public, canonical resources that AI assistants and crawlers should use to understand this site.

Site Overview

- Canonical URL: <https://newsletter.surf/>
- Site type: web site
- Recommended summary: A curated directory of newsletters worth your inbox. Browse hand-reviewed newsletters by category - tech, business, culture, and more.

Core URLs

- [Homepage](<https://newsletter.surf/>): Primary public entry point and canonical site overview.
 - [Search](<https://newsletter.surf/search>): Important public page discovered from the homepage navigation.
 - [Tools](<https://newsletter.surf/tools>): Important public page discovered from the homepage navigation.
 - [Guides](<https://newsletter.surf/guides>): Documentation and implementation guidance for the site.
- Continued in the full scan report...

robots.txt additions

TXT

```
# robots.txt additions
# Copy these blocks into the existing robots.txt file. Keep current rules
unless a note calls out a conflicting Disallow.

# AI crawler access
# Add explicit Allow rules for blocked AI crawlers; remove or narrow
conflicting Disallow rules if your crawler target requires precedence.
User-agent: GPTBot
Allow: /

User-agent: ChatGPT-User
Allow: /

User-agent: ClaudeBot
Allow: /

User-agent: Claude-Web
Allow: /

User-agent: PerplexityBot
Allow: /

User-agent: Google-Extended
Allow: /

# Crawl-rate guidance
Continued in the full scan report...
```

schema.json

JSON

```
{
"@context": "https://schema.org",
"@graph": [
{
"@type": "Organization",
"@id": "https://newsletter.surf/#organization",
"name": "Newsletter Surf - Discover great newsletters",
"description": "A curated directory of newsletters worth your inbox. Browse
hand-reviewed newsletters by category - tech, business, culture, and more.",
"url": "https://newsletter.surf/",
"logo":
"https://cdn.newsletter.surf/newsletters/newsletters-2973962c-2470-4f41-a415
-914bdf3b5e8c.png"
},
{
"@type": "WebSite",
"@id": "https://newsletter.surf/#website",
"name": "Newsletter Surf - Discover great newsletters",
"description": "A curated directory of newsletters worth your inbox. Browse
hand-reviewed newsletters by category - tech, business, culture, and more.",
"url": "https://newsletter.surf/",
"publisher": {
"@id": "https://newsletter.surf/#organization"
},
"inLanguage": "en",
"potentialAction": {
Continued in the full scan report...
```

Content-Signal

TXT

Content-Signal recommendations

Use these directives to make AI-use preferences explicit for compliant crawlers and AI systems. They are advisory signals, so keep them aligned with robots.txt, terms, and access controls.

Recommended values

- ai-train=no: AI model training, fine-tuning, and dataset creation.
- search=yes: AI search indexing, snippets, and discovery.
- ai-input=yes: AI answer grounding, retrieval, and generated-response context.

HTTP response header

```
'''http
Content-Signal: ai-train=no, search=yes, ai-input=yes
Content-Usage: train-ai=n
'''
```

Best for site-wide or route-specific policies because the signal travels with every response, including pages that AI systems fetch directly.

HTML meta tag alternatives

Add these inside the document '`<head>`' when server headers are not

continued in the full scan report...

Full report

<https://llmscan.dev/scan/QVfNsibz59uPyBtDvBNhl>