

# axray.online

<https://axray.online/>

Scanned 2026-09-25 20:02 UTC against AX spec v1.5.7

100 / 100  
Grade A+

|                                             |             |     |
|---------------------------------------------|-------------|-----|
| Reachability<br>25% weight - nothing lost   | <div></div> | 100 |
| Comprehension<br>25% weight - nothing lost  | <div></div> | 100 |
| Structure<br>18% weight - nothing lost      | <div></div> | 100 |
| Actionability<br>17% weight - nothing lost  | <div></div> | 100 |
| Agent Contract<br>15% weight - nothing lost | <div></div> | 100 |

IN PLAIN ENGLISH

**An AI assistant can find axray.online, read it, and act on what it finds.**

**Can an AI assistant open this page at all?**

**YES** Yes. When an assistant goes to look, the page answers it.  
This is the thing most sites get wrong first, and this one does not.

**Will it describe this business correctly?**

**YES** Yes. The page states plainly what it is, in words the assistant can read.  
What an assistant says about this page comes from the page itself.

**Can it quote the facts - prices, contact details, opening hours?**

**YES** Yes. The key facts are published in a form it can lift word for word.  
Asked what this costs or how to get in touch, an assistant answers with the real number.

**Can it take the next step - follow a link, use a form, get in touch?**

**YES** Yes. Links lead somewhere, forms are labelled, and there is a clear way to make contact.  
An assistant can carry somebody from an answer all the way to you.

## Have the AI companies been told your rules?

YES

Yes. The site states which AI systems may use it and on what terms.  
The terms are yours, in writing, where the crawlers look for them.

---

WHAT CHANGES IF YOU ACT

## This page is already agent-ready

Every question above is already answered "yes". The work now is keeping it that way: run the scan in your deploy pipeline so a change that breaks it fails the build rather than being discovered months later.

## Nothing is failing

No check on this page is failing or warning. The passing list below is the whole report.

---

64 PASSING

### What this page already gets right

- PASS Page responds to a plain HTTP request
- PASS No bot wall in front of the content
- PASS Content exists without running JavaScript
- PASS The status code matches what the page says
- PASS robots.txt is present and parseable
- PASS Answer engines are allowed to crawl
- PASS User-triggered fetchers are allowed
- PASS Autonomous agents are allowed
- PASS Retrieval providers are allowed
- PASS Classic search crawlers are allowed
- PASS No directive suppressing AI use of this page
- PASS Short redirect chain
- PASS Responds fast enough for an agent budget
- PASS HTML payload stays inside an agent context budget
- PASS HTML is compressed on the wire
- PASS Re-fetching is cheap for a crawler
- PASS Declares a correct content type
- PASS Has a descriptive <title>
- PASS Has a meta description
- PASS The page describes itself consistently
- PASS Exactly one <h1> naming the page
- PASS Headings form a sane outline
- PASS Enough substance to be worth citing
- PASS Answers questions in a form that can be quoted
- PASS Passages make sense on their own
- PASS Signal-to-markup ratio
- PASS Declares a document language
- PASS Text arrived intact, not mangled
- PASS Images carry alt text
- PASS Main content is marked as such
- PASS Uses semantic elements, not div soup
- PASS URL describes the content
- PASS Ships structured data
- PASS Structured data parses cleanly
- PASS Structured data describes the right thing

|      |                                                |
|------|------------------------------------------------|
| PASS | Structured data agrees with the page           |
| PASS | Open Graph card is complete                    |
| PASS | Declares a canonical URL                       |
| PASS | Publishes a sitemap                            |
| PASS | Publishes a feed of its content                |
| PASS | Declares its language alternates               |
| PASS | Prices are machine-readable                    |
| PASS | Dates are machine-readable                     |
| PASS | Declares a link preview card                   |
| PASS | Navigation uses real links                     |
| PASS | Link text describes the destination            |
| PASS | Links onward into the site                     |
| PASS | Form fields are labelled                       |
| PASS | Form fields have stable names                  |
| PASS | Fields declare autocomplete tokens             |
| PASS | Controls are real buttons                      |
| PASS | Reachable by a human, discoverably             |
| PASS | Site search is machine-usable                  |
| PASS | Declares what an agent can do here             |
| PASS | Programmatic surface is discoverable           |
| PASS | Publishes an llms.txt                          |
| PASS | llms.txt is actually useful                    |
| PASS | States an explicit AI policy                   |
| PASS | Declares AI usage preferences the standard way |
| PASS | Exposes an agent-callable manifest             |
| PASS | Identity is machine-verifiable                 |
| PASS | Says when the content changed                  |
| PASS | Declares content licensing                     |
| PASS | Publishes security.txt                         |

6 checks did not apply to this page and were excluded from the score entirely: Coding agents can read the documentation; Real agent user-agents are served normally; Provides a no-JavaScript fallback; Tabular data is marked up as a table; Declares its place in the site; Inline microdata annotations.

## Who can actually reach this page

Resolved from robots.txt for this exact path, one crawler at a time. Blocking a training crawler is a legitimate business decision. Blocking an answer engine removes you from the results that would have linked to you.

### Answer engine

Indexes you so an assistant can cite and link you. Blocking it removes you from AI answers.

|                  |            |         |
|------------------|------------|---------|
| OAI-SearchBot    | OpenAI     | Allowed |
| Claude-SearchBot | Anthropic  | Allowed |
| PerplexityBot    | Perplexity | Allowed |
| Applebot         | Apple      | Allowed |
| meta-webindexer  | Meta       | Allowed |
| Amazonbot        | Amazon     | Allowed |
| DuckAssistBot    | DuckDuckGo | Allowed |
| Bravebot         | Brave      | Allowed |
| YouBot           | You.com    | Allowed |
| TongyiBot        | Alibaba    | Allowed |
| PetalBot         | Huawei     | Allowed |
| YandexAdditional | Yandex     | Allowed |

### User-triggered

Fetches your page because a human asked an assistant to open it right now. Blocking it breaks a live request.

|                        |             |         |
|------------------------|-------------|---------|
| ChatGPT-User           | OpenAI      | Allowed |
| Claude-User            | Anthropic   | Allowed |
| Perplexity-User        | Perplexity  | Allowed |
| MistralAI-User         | Mistral     | Allowed |
| Google-NotebookLM      | Google      | Allowed |
| GoogleAgent-URLContext | Google      | Allowed |
| kagi-fetcher           | Kagi        | Allowed |
| meta-externalfetcher   | Meta        | Allowed |
| Kimi-User              | Moonshot AI | Allowed |
| Amzn-User              | Amazon      | Allowed |
| cohere-ai              | Cohere      | Allowed |

### Autonomous agent

Browses, compares and buys on a person's behalf. Blocking it removes you from the shortlist before anyone sees it.

|                      |                  |         |
|----------------------|------------------|---------|
| ChatGPT Agent        | OpenAI           | Allowed |
| Operator             | OpenAI           | Allowed |
| GoogleAgent-Mariner  | Google           | Allowed |
| Gemini-Deep-Research | Google           | Allowed |
| AmazonBuyForMe       | Amazon           | Allowed |
| NovaAct              | Amazon           | Allowed |
| Manus-User           | Butterfly Effect | Allowed |

### Retrieval provider

Sells the index that other companies' AI products search. Blocking it removes you from all of them at once.

|              |        |         |
|--------------|--------|---------|
| ExaSearchBot | Exa    | Allowed |
| TavilyBot    | Tavily | Allowed |

|                    |            |         |
|--------------------|------------|---------|
| FirecrawlAgent     | Firecrawl  | Allowed |
| ShapBot            | Parallel   | Allowed |
| Diffbot            | Diffbot    | Allowed |
| Cloudflare-AutoRAG | Cloudflare | Allowed |
| bedrockbot         | Amazon     | Allowed |

## Classic search

Traditional search indexing, increasingly the substrate AI answers are built on.

|           |           |         |
|-----------|-----------|---------|
| Googlebot | Google    | Allowed |
| Bingbot   | Microsoft | Allowed |

## Coding agent

Reads documentation while a developer is building against you. Blocking it is how your API gets used wrong.

|                   |           |         |
|-------------------|-----------|---------|
| Claude-Code       | Anthropic | Allowed |
| Cursor            | Cursor    | Allowed |
| Google-Gemini-CLI | Google    | Allowed |
| Devin             | Cognition | Allowed |
| opencode          | opencode  | Allowed |
| Trae              | ByteDance | Allowed |

## Training crawler

Collects content for model training. Blocking it is a legitimate business choice with no traffic cost.

|                    |              |         |
|--------------------|--------------|---------|
| GPTBot             | OpenAI       | Allowed |
| ClaudeBot          | Anthropic    | Allowed |
| Google-Extended    | Google       | Allowed |
| Applebot-Extended  | Apple        | Allowed |
| meta-externalagent | Meta         | Allowed |
| CCBot              | Common Crawl | Allowed |
| Bytespider         | ByteDance    | Allowed |
| DeepSeekBot        | DeepSeek     | Allowed |
| AI2Bot             | AI2          | Allowed |

---

WHAT SURVIVES BEING CUT UP

## The pieces an assistant is actually handed

No assistant reads a whole page. A retrieval system cuts it into passages of a few hundred tokens and hands back one of them, and the answer is written from that passage alone. This page cuts into 23 passages of about 300 tokens; 100% of its own content sits in a passage that can be read on its own, and 0% of what would be retrieved is navigation rather than content.

**Passage 1**      no heading above it

It carries a byline, an image or a bare link and no words of its own, so retrieved on its own it says nothing.

---

## What we actually received

Every number in this report comes from one plain HTTP response. No browser, no JavaScript, no retries.

|                |                                                        |
|----------------|--------------------------------------------------------|
| Final URL      | <code>https://axray.online/</code>                     |
| HTTP status    | 200                                                    |
| Response time  | 200 ms                                                 |
| HTML size      | 53 KB                                                  |
| Inline script  | 5 KB (9% of payload)                                   |
| Readable words | 1792                                                   |
| Text-to-markup | 19.3%                                                  |
| Title          | AXRAY - See what AI can actually read on your website. |
| Language       | en                                                     |
| Canonical      | <code>https://axray.online/</code>                     |
| Detected stack | (none detected)                                        |
| robots.txt     | served                                                 |
| sitemap.xml    | 241 URLs                                               |
| llms.txt       | 11 sections, 41 links                                  |

---

Generated by AXRAY. The live version of this report, with every check expandable, is at <https://axray.online/report/gv7UEt8pKqGqcA> - re-run it any time to see whether a change worked.