

MENTIONSHARE · GEO CONSULTING

Attention Insight

AI Visibility Audit and Fix Plan

Engines tested: ChatGPT, Claude, Perplexity, Gemini, Google AI Overviews

Test date: 2026-08-16 · Locale: Lithuania (English queries)

Prepared by MentionShare

Checked by Laurynas Leskauskas

Executive summary

62 / 100

Provisional score, Developing band. Visibility Now 37/55. Readiness 25/45. This page

and the program view are all you need to act; every fix card runs on its own, and the appendix exists to prove the claims.

Scope note: all 5 engine panels ran in full on your 5 buyer questions. The scoring-weights document is pending, so the numbers above are provisional and will be recomputed without a re-test. Several review platforms block automated readers; where their figures appear, the report says so and treats them as claims from their public listings, not verified fact.

The three most damaging findings.

- **Wrong prices sit where AI reads.** Your [Capterra listing](#) still shows a retired plan structure: a Solo plan at \$23 with a 7-day trial and a Team plan at \$479. Those plans no longer exist. Claude read that listing and quoted the whole structure to a buyer as your current pricing (ACC-02, SUR-03). A second wrong figure, "from about \$19/month", circulates on Perplexity (ACC-03). Fixed by [FIX-01](#) and [FIX-10](#).
- **You lose the alternatives question.** When buyers ask for Neurons Predict alternatives, the highest-intent question in your panel, 3 of 5 engines never say your name. They quote roundup lists and competitor pages where you are absent, including a roundup from 2019 (Q3, R13 to R15, GAP-01, CIT-15). You rank first on 10 of the 20 generic answers, so the gap is this question, and it is fixable. Fixed by [FIX-04](#) and [FIX-08](#).
- **Your money page splits its identity.** Your pricing lives at two addresses at once: /pricing/ tells machines its true copy is /billing-plans/, the address buyers and engines actually cite is missing from your sitemap, and the page carries no machine-readable price data at all (SUR-16, SUR-02, SUR-19). Engines split one page's credit across two records, and the correct prices are harder to lift than they should be. Fixed by [FIX-02](#) and [FIX-09](#).

Your first three moves.

- [FIX-01](#) Correct your Capterra pricing listing. Owner: Marketing. Quick win.
- [FIX-02](#) Give your pricing page one address and put it in the sitemap. Owner: Developer. Quick win.
- [FIX-02](#) Approve one accuracy sentence and deploy it everywhere. Owner: Marketing. Quick win.

90-day outlook. If the first three moves ship this week, the corrected Capterra listing and the unified pricing page reach the surfaces engines re-read on every answer. Fixes on live-retrieved surfaces (your pages, your profiles, articles that cite you) can show up in answers within days to weeks of a recrawl. Answers rooted in model training move on model-update timescales that nobody outside the AI labs controls. Winning the alternatives question depends on the comparison content and roundup work, which compounds more slowly.

Scorecard

Provisional: recomputed without re-testing once the pending weights land.

Component	Points	What this means for you
AI share of voice	18 / 25	You are named in 20 of 25 answers and ranked first in half of the generic ones; the misses cluster on one question.
Accuracy of AI claims	6 / 10	Engines reading your live site quote you correctly; engines reading Capterra or price aggregators quote plans and figures you retired.
How AI frames you	7 / 10	Engines call you the best value and the default pick; ChatGPT alone consistently places a competitor above you for professional depth.
Crawler access	9 / 10	Your robots.txt (the small public file that tells bots what they may read) allows every AI crawler, and your copy is fully readable without JavaScript.
Structured data	5 / 10	Your pages carry general machine-readable labels, but your pricing page carries none at all: no product, no prices, nothing an engine can lift as data.
Page structure	4 / 10	No page opens with an answer, no heading matches a buyer question except accuracy, and your whole site contains zero real HTML tables.
Positioning clarity	3 / 5	Your category naming is consistent; your accuracy number circulates in four versions, and sources disagree on your founding year.
Content coverage	4 / 10	The questions you lose have no page on your site: no alternatives comparison, no category cost guide. Your case studies are a real strength.
Off-site presence	6 / 10	You are present and rated 4.8 on the review platforms buyers check, but your Capterra facts are stale and your Trustpilot has been quiet for about a year.

Overall: 62 / 100, provisional. Plain reading: the market already likes you, and the engines mostly recommend you; the risks are wrong third-party facts, one lost question, and technical housekeeping on the money page. Sub-scores: Visibility Now 37/55, Readiness 25/45.

Program view

This page plus the cards is your operating plan. Work top to bottom. A card that needs another card finished first says so at the bottom; no line there means the card can start any time.

Fix	Title	Owner	Effort	Lane
This week				
FIX-01	Correct your Capterra pricing listing	Marketing	Quick win	Data hygiene
FIX-02	Give your pricing page one address and put it in the sitemap	Developer	Quick win	Technical access
FIX-03	Approve one accuracy sentence and deploy it everywhere	Marketing	Quick win	Data hygiene
This month				
FIX-04	Publish your own Neurons Predict alternatives comparison	Marketing	Medium	Content
FIX-05	Publish the predictive eye tracking cost guide	Marketing	Medium	Content
FIX-06	Settle your founding year and align the company databases	Founder	Quick win	Data hygiene
FIX-09	Add price data machines can read to the pricing page	Developer	Medium	Technical access
FIX-10	Sweep the software directories for retired facts	Marketing	Medium	Data hygiene
This quarter				
FIX-07	Restart the review flow on Trustpilot and G2 (recurring)	Marketing	Quick win	Off-site and links
FIX-08	Get into the alternatives roundups engines actually quote	Marketing	Project	Off-site and links
FIX-11	Publish llms.txt	Developer	Quick win	Technical access

Fix plan

FIX-01 · Correct your Capterra pricing listing

Data hygiene · Impact: High · Effort: Quick win · Owner: Marketing

The problem. [Capterra](#) still shows pricing plans you retired, and AI tools read Capterra and repeat them to buyers.

Why this matters to you. Claude answered a buyer question about you by quoting a Solo plan at \$23 with a 7-day trial and a Team plan at \$479, all read from Capterra, and none of it exists in your current lineup (ACC-02, SUR-03). A buyer comparing you on those numbers is comparing a product you stopped selling. Correcting the listing corrects the answer at its source.

Do this.

1. Marketing signs into the Gartner Digital Markets vendor portal (digitalmarkets.gartner.com), which manages Capterra listings; if the product profile is unclaimed, start at capterra.com/vendors/. You will see the current listing values while editing, which is the point where the stale data becomes visible to you.
2. Replace the pricing section with the corrections table below, matching your live pricing page exactly.
3. Check the sister listings that share Capterra's data platform, [GetApp](#) and [Software Advice](#), from the same portal and apply the same corrections.
4. Seven days later, open your [public Capterra page](#) logged out and confirm the new plans and the 14-day trial show.

Ready to use. Corrections for the listing:

Field	Listing shows	Correct value (your live pricing page)
Entry plan	Solo, \$23/month, 1 seat, 20 credits, watermark, 3 months storage	Basic, EUR 29/month (\$31), 1 seat, 40 credits, watermark
Mid plan	not shown	Pro, EUR 119/month (\$129), 2 seats, 200 credits, no watermark, API
Top self-serve plan	Team, \$479/month, 5 seats, 500 credits, 12 months storage	Hero, EUR 299/month (\$324), 2 seats, unlimited credits fair use, SSO
Free trial	7-day	14 days, no credit card required
Annual billing	not shown	20% off: Basic EUR 23, Pro EUR 95, Hero EUR 250 per month

Done when. The public Capterra listing, viewed logged out, shows Basic/Pro/Hero and the 14-day trial.

FIX-02 · Give your pricing page one address and put it in the sitemap

Technical access · Impact: High · Effort: Quick win · Owner: Developer

The problem. Your pricing lives at two addresses at once, the address buyers know is missing from your sitemap, and your Figma plugin page is split across three addresses.

Why this matters to you. /pricing/ tells machines its true copy is /billing-plans/, /billing-plans/ is what sits in your sitemap, and both serve the same content (SUR-16, SUR-02). Engines dilute one page's authority across two records, and the address every third-party site links to is the one your own sitemap omits. One address means one record that collects all the credit.

Do this.

1. Developer picks /pricing/ as the single pricing address (it is the address engines already cite in answers).
2. In WordPress, set the canonical of /pricing/ to itself, and 301-redirect /billing-plans/ to /pricing/ (redirect rule below).
3. Consolidate the Figma plugin the same way: keep /attention-insight-figma-plugin/ (it is the one in your sitemap), 301 /figma-plugin-101/ and /figma-plugin/ to it.
4. Confirm /pricing/ now appears in page-sitemap.xml, then resubmit the sitemap index in Google Search Console and Bing Webmaster Tools.
5. Check the live result: fetch /billing-plans/ and confirm it lands on /pricing/ with one hop.

Ready to use. Redirect and canonical rules:

```
301: https://attentioninsight.com/billing-plans/ -> https://attentioninsight.com/pricing/
301: https://attentioninsight.com/figma-plugin-101/ -> https://attentioninsight.com/attention-
insight-figma-plugin/
301: https://attentioninsight.com/figma-plugin/ -> https://attentioninsight.com/attention-insight-
figma-plugin/

canonical on /pricing/: <link rel="canonical" href="https://attentioninsight.com/pricing/" />
```

Done when. /pricing/ self-canonicals, sits in the sitemap, and both old addresses 301 to their keepers.

FIX-03 · Approve one accuracy sentence and deploy it everywhere

Data hygiene · Impact: High · Effort: Quick win · Owner: Marketing

The problem. Your accuracy number, the reason people buy, circulates in at least four versions, and one engine has started inventing its own.

Why this matters to you. In plain terms: the panel heard 90%, 92.5%, 94.5%, and 96%, and Gemini told a buyer you trained on "over 70,000 eye-tracking studies", a figure no source carries (ACC-04, ACC-10, CIR-03). Engines trust facts that agree across sources. One approved sentence, deployed everywhere, ends the guessing.

Do this.

1. Marketing approves the drafted sentence below or adjusts its numbers to the ones you stand behind.
2. Place it verbatim on /technology/ directly under the "Accuracy validation" heading, as the first sentence.
3. Deploy it through the checklist below, one surface at a time, replacing any older accuracy wording.
4. Confirm the live result: search your own site for "92" and "96" and check every hit now matches the approved sentence.

Ready to use. The drafted canonical sentence:

Attention Insight's attention predictions reach up to 96% accuracy compared with real eye-tracking studies, scoring 92.5% on the MIT saliency benchmark, with a model trained on more than 5.5 million recorded eye-tracking fixations.

Deployment checklist: attentioninsight.com/technology/ · homepage accuracy mention · Capterra listing description (with [FIX-01](#)) · G2 product description · Chrome Web Store listing text · Figma plugin listing text · press kit page.

Done when. The same sentence, word for word, appears on your site and in your two biggest review-platform descriptions.

FIX-04 · Publish your own Neurons Predict alternatives comparison

Content · Impact: High · Effort: Medium · Owner: Marketing

The problem. When buyers ask AI for alternatives to your biggest competitor, three engines answer from lists you are not on, and your site has no page in the conversation.

Why this matters to you. The alternatives question is where buyers compare before choosing, and you are absent from 3 of 5 answers while ranked first on most other questions (Q3, R13 to R15, GAP-01). Engines lift the passage whose heading matches the query, and today no heading on your site matches it. The page below gives them one to lift.

Do this.

1. Marketing commissions the brief below to your writer; target publish inside 3 weeks at /neurons-predict-alternatives/.
2. Week 1: writer completes the text from the brief; you verify every competitor fact against the named public page before approval, and keep a checked-facts sheet next to the article.

3. Week 2: publish, add the page to your main navigation under Resources, and link to it from /features/ and the pricing page with descriptive anchors.
4. Week 3: submit the URL in Google Search Console and Bing Webmaster Tools, then re-ask the exact panel question in ChatGPT and Perplexity and record whether the page is cited; repeat the check after 30 days.

Ready to use. Content brief. Working title: Neurons Predict alternatives: 5 predictive attention tools compared. Target questions, used verbatim as H2 headings: "Neurons Predict alternatives?" and "best predictive eye tracking software for UX designers?" (Q2, Q3). Drafted opening (place first, 56 words):

Looking for Neurons Predict alternatives? The five predictive attention tools most teams compare are Attention Insight, Brainsight, EyeQuant, VisualEyes, and expoze.io. All five predict where people look without recruiting participants; they differ on price transparency, integrations, and validation. This comparison covers each tool's entry price, accuracy validation, and design-workflow fit, with sources you can check.

Comparison table skeleton, rows the audit can already fill: tool · entry price (Attention Insight EUR 29/mo; others from their public pages at write time) · free trial (yours: 14 days, no card) · validation basis (yours: MIT saliency benchmark, up to 96%) · integrations (yours: Figma, Chrome, Adobe Express, Photoshop, InDesign, Sketch) · watermark-free tier. Statistic slots: your accuracy sentence from [FIX-Q3](#); one dated case-study result (Filmelior, 258% CTA clicks, from /case-studies/). Internal links: /pricing/, /technology/, /case-studies/. Honest-comparison rule: every competitor's genuine strength stated, every claim traced to a public page with the date checked. Length: 1,600 to 2,200 words.

Done when. The page is live with the two H2s verbatim, linked from the navigation, and indexed in both consoles.

FIX-05 · Publish the predictive eye tracking cost guide

Content · Impact: High · Effort: Medium · Owner: Marketing

The problem. Buyers asking what this category costs get answers stitched from aggregators, one of which prices you at \$19, and your site has no page that answers the question.

Why this matters to you. Claude's cost answer skipped you entirely, and Perplexity priced you at a figure that matches no plan you sell (R17, ACC-03, GAP-04). The brand that owns the category cost answer gets named when the buying question follows, and your real prices are a selling point you currently let third parties misquote.

Do this.

1. Marketing commissions the brief below; target publish inside 3 weeks at /predictive-eye-tracking-cost/.
2. Week 1: writer completes the text; verify every third-party price against the vendor's live page on writing day and date every figure in the text.
3. Week 2: publish with the price table as a real HTML table, add internal links from /pricing/ and the [FIX-Q4](#) comparison page.
4. Week 3: submit the URL in both webmaster consoles; re-ask the panel cost question in Claude and Perplexity after 30 days and record whether the page is cited.

Ready to use. Content brief. Working title: How much does predictive eye tracking software cost in 2026? Target questions, used verbatim as H2 headings: "how much does predictive eye tracking software cost?"

What is the best service?" (Q4) and "what is the best attention heatmap tool to test designs before launch?" (Q1). Drafted opening (place first, 54 words):

Predictive eye tracking software costs between roughly EUR 20 and EUR 300 per month for self-serve plans, with enterprise tiers priced by quote. Attention Insight starts at EUR 29 per month with a 14-day free trial. This guide compares what the entry price of each major tool includes, verified against live pricing pages.

Price table spec (real HTML table): plan · monthly price · annual price · credits or analyses included · seats · trial. Your row uses live values: Basic EUR 29 (\$31), Pro EUR 119 (\$129), Hero EUR 299 (\$324), annual EUR 23/95/250, 14-day trial, credit definition (1 credit = 1 image or URL analysis). Statistic slots: your accuracy sentence ([FIX-03](#)); one external dated category statistic if a checkable one exists at write time, otherwise omit. Internal links: /pricing/, /neurons-predict-alternatives/, /case-studies/. Length: 1,400 to 2,000 words.

Done when. The page is live with the Q4 wording as an H2, a real HTML price table with your current values, and indexed in both consoles.

FIX-06 · Settle your founding year and align the company databases

Data hygiene · Impact: Medium · Effort: Quick win · Owner: Founder

The problem. Public records disagree about when your company was founded, and engines repeat whichever record they happen to read.

Why this matters to you. Crunchbase says founded 2019; press coverage says 2018; Claude told a buyer 2019 while citing Crunchbase (ACC-07, SUR-06, SUR-07). Contradictions between sources make engines hedge or pick wrong, and identity facts are the easiest ones to make agree.

Do this.

1. Founder confirms the three facts in the Needs-from-you list below, in one email.
2. Marketing updates Crunchbase first (you control the company profile after claiming it), then your own about page so both carry the same year and office list.
3. Where a database you cannot edit carries the wrong year, use its listed correction contact; you will see each listing's current values as you open them, which is where the remaining conflicts become visible.

Needs from you. 1) The founding year you stand behind (2018 or 2019). 2) The current office list (Vilnius plus which others). 3) Current headcount for database profiles.

Done when. Crunchbase and your own about page state the same founding year and offices.

FIX-07 · Restart the review flow on Trustpilot and G2 (recurring)

Off-site and links · Impact: Medium · Effort: Quick win · Owner: Marketing · (recurring)

The problem. Your Trustpilot page has been quiet for about a year, and engines read review recency as a signal of whether a product is alive.

Why this matters to you. ChatGPT told a buyer your Trustpilot shows no new reviews in 12 months (ACC-08, SUR-05). Your ratings are excellent, 4.8 on G2 and Capterra per their listings, so the fix is volume and recency, and reviews are quoted verbatim in AI answers.

Do this.

1. Marketing owns this card and sets a monthly calendar reminder: first Monday, send the drafted request below to the 5 customers with the best recent results.
2. Rotate the ask: odd months [Trustpilot](#), even months [G2](#), so both surfaces stay current.
3. Log outcomes in a sheet (sent, review posted, no response) and follow up once after 7 days.
4. Close condition each quarter: at least 3 new reviews across the two platforms; if a quarter misses, raise the send list to 10.

Ready to use. Drafted request (94 words):

Subject: A two-minute favor

Hi [name],

You recently used Attention Insight on [project], and the results you shared with us were exactly what the tool is for. Would you be willing to spend two minutes leaving an honest review on [Trustpilot/G2]? Here is the direct link: [link].

One or two sentences about what you tested and what you found is perfect. Honest reviews, including the critical parts, help other teams decide if the tool fits.

Thanks either way,
[your name]

Done when. The monthly reminder exists in the owner's calendar and the first batch of 5 requests is sent.

FIX-08 · Get into the alternatives roundups engines actually quote

Off-site and links · Impact: High · Effort: Project · Owner: Marketing

The problem. The lists engines quote when buyers ask for alternatives mostly leave you out, including one written in 2019.

Why this matters to you. In plain terms: when Perplexity, Gemini, and Google AI Overviews answered the alternatives question, they drew from roundup articles and competitor pages where you do not appear (R13 to R15, CIT-12 to CIT-16). Your own comparison page ([FIX-04](#)) is half the answer; the other half is appearing in the lists those engines already trust.

Expectations, plainly. This is a numbers game: a short email is worth a try, not a promised win, and most pitches get no reply. Some targets only add vendors for payment; a paid placement is an advertising decision outside this fix plan, so log "asked for payment" as an outcome and move on. A good result looks

like one addition from every eight attempts. Direct submissions need no reply from anyone, so they run first and the card produces something even if every email is ignored.

Do this.

1. Marketing copies the target table below into a tracking sheet and works it top to bottom, direct submissions first.
2. For article targets, find the person: open the article, take the byline author, and find their contact through the byline link, the site's contact page, or LinkedIn; never send to a generic inbox when a byline exists.
3. Send the drafted pitch below, one concrete give included; follow up once after 10 days, then record the outcome and stop.
4. Weekly cadence for 6 weeks, then close the card by tallying outcomes; re-open next quarter with fresh targets found by re-asking the panel question.

Ready to use. Target list (statuses are yours to fill in as you work the sheet):

Target	Why this one	How to get in	Status today (you fill this in at step 1)
Capterra "Attention Insight alternatives" category page	Capterra data feeds several engine answers in your panel	Direct: vendor portal category and alternatives settings	
G2 predictive-attention category and comparison pages	Cited on reputation answers; comparisons are indexed	Direct: G2 vendor portal	
AlternativeTo listing for Neurons and for Attention Insight	Alternatives queries lean on AlternativeTo-style lists	Direct submission; community-moderated, entry can be edited by users, completeness keeps it accurate	
The 2019 Medium roundup Google AI Overviews cited (Paul Boag, UX design tools)	Still feeding answers 7 years on (CIT-15)	Email the byline author with an update note	
2026 "best predictive eye tracking" roundups found by searching the exact panel question	Perplexity and Gemini cite "2026 roundups" generically (CIT-16)	Email the byline author per article	

Drafted pitch (118 words, one concrete give):

Subject: An update for your predictive eye tracking roundup

Hi [name],

Your piece on [article title] still ranks for [query], and AI assistants quote it when people ask for design-testing tools. One update you might find useful: Attention Insight now validates at up to 96% against the MIT saliency benchmark, starts at EUR 29 a month with a 14-day trial, and ships Figma, Chrome, and Adobe integrations.

If it helps your readers, we can share our internal comparison table of the five main predictive attention tools, prices and validation methods checked this month, free to reuse with or without credit.

Either way, thanks for the piece.

[your name], Attention Insight

Done when. All targets in the sheet carry an outcome (added, declined, asked for payment, no reply after follow-up), with any number of additions counted as the result.

Depends on: [FIX-04](#)

FIX-09 · Add price data machines can read to the pricing page

Technical access · Impact: Medium · Effort: Medium · Owner: Developer

The problem. Your pricing page carries no machine-readable data at all, and no page on your site contains a single real HTML table.

Why this matters to you. In plain terms: engines lift prices most reliably from labeled data and real tables, and your pricing page offers neither (SUR-19, EX-05). Your prices survived in answers because your raw text is readable; labeled data makes the correct numbers the easiest thing to quote and protects you from the wrong ones (ACC-02, ACC-03).

Do this.

1. Developer adds the JSON-LD block below (JSON-LD is a small script tag with labeled facts machines parse) to /pricing/, updating values whenever pricing changes.
2. Rebuild the plan grid so the plan-versus-feature comparison also renders as a real HTML table in the page source, not only styled boxes.
3. Validate the block at validator.schema.org and fix any warning before publishing.
4. Check the live result: view the page source and confirm the script tag and the table element are present.

Ready to use. Filled JSON-LD (validated structure, your live values):

```
<script type="application/ld+json">
{
  "@context": "https://schema.org",
  "@type": "Product",
  "name": "Attention Insight",
  "description": "AI-powered attention analytics that predicts where people look on designs, ads, and websites using predictive eye-tracking heatmaps.",
  "brand": { "@type": "Brand", "name": "Attention Insight" },
  "url": "https://attentioninsight.com/pricing/",
  "offers": [
    { "@type": "Offer", "name": "Basic", "price": "29", "priceCurrency": "EUR",
      "description": "1 user seat, 40 credits per month, watermark on outputs" },
    { "@type": "Offer", "name": "Pro", "price": "119", "priceCurrency": "EUR",
      "description": "2 user seats, 200 credits per month, no watermark, API access" },
    { "@type": "Offer", "name": "Hero", "price": "299", "priceCurrency": "EUR",
      "description": "2 user seats, unlimited credits under fair use, SSO, expert feedback" }
  ]
}
</script>
```

Done when. validator.schema.org passes the block with zero errors and the plan comparison exists as a table element in the page source.

Depends on: [FIX-02](#)

FIX-10 · Sweep the software directories for retired facts

Data hygiene · Impact: Medium · Effort: Medium · Owner: Marketing

The problem. Beyond Capterra, several directories and aggregators carry fragments of your old pricing, and at least one engine already quoted a price you never charged.

Why this matters to you. Perplexity's "from about \$19/month" matches no plan you sell, and its likely home is a stale directory or coupon aggregator (ACC-03). Engines read these surfaces constantly, so every stale row is a wrong answer waiting to be generated.

Do this.

1. Marketing works the table below top to bottom; for each surface, open your listing and compare its pricing row against your live pricing page; you will see each listing's current values during the visit, which is where the wrong ones surface.
2. Correct what you control through each platform's vendor or edit flow; for surfaces without vendor access, use their listed correction contact.
3. Log per surface: checked date, what it showed, corrected or contact-sent.
4. Re-check the corrected surfaces after 30 days; aggregators re-sync slowly.

Ready to use. Sweep list (what each showed goes in your log as you open them): [GetApp](#) · [Software Advice](#) · [SaaSwothy](#) · SourceForge · AlternativeTo · [tenereteam](#) and similar coupon scrapers (request removal or correction of pricing rows) · [AppSumo](#) listing text (the deal era is over; ask for the listing description to state current pricing model). For each: field to check is entry price, plan names, trial length, against Basic EUR 29 / Pro EUR 119 / Hero EUR 299, 14-day trial.

Done when. Every surface in the list carries a logged outcome and no checked listing shows a \$19 or Solo/Team row.

FIX-11 · Publish llms.txt

Technical access · Impact: Low · Effort: Quick win · Owner: Developer

The problem. Your site has no llms.txt, a simple text file some companies now publish as a summary written for AI tools.

Why this matters to you. A note on llms.txt: adoption is growing but no major AI engine has committed to reading the file, and current data shows no measurable citation benefit from having one. We include it because it costs minutes, carries no risk, and doubles as a clean machine-readable summary of your company.

Do this.

1. Developer saves the file below to your web root as /llms.txt.
2. Confirm it resolves at <https://attentioninsight.com/llms.txt> as plain text.
3. Re-check it each quarter alongside the [FIX-07](#) cadence, updating prices if they changed.

Ready to use.

```
# Attention Insight
> Attention Insight is an AI-powered attention analytics platform that predicts where
> people look on designs, ads, and websites using predictive eye-tracking heatmaps, so
> teams can test and fix designs before launch without recruiting participants.

## Product
- [Features](https://attentioninsight.com/features/): Heatmaps, Clarity Score, Focus Score, AI
  recommendations
- [Technology](https://attentioninsight.com/technology/): Model, dataset, and accuracy validation (up
  to 96% vs eye-tracking studies)
- [Pricing](https://attentioninsight.com/pricing/): Plans from EUR 29 per month, 14-day free trial

## Proof
- [Case studies](https://attentioninsight.com/case-studies/): Documented results with numbers

## Company
- [About](https://attentioninsight.com/): Team, founding, and contact
```

Done when. The file resolves at /llms.txt and matches your canonical description word for word.

Evidence appendix

Every ID cited above resolves to a row here. Compressed on purpose; the cards carry the plan.

E1. Where you stand

25 answers (5 questions x 5 engines). Named once per answer, exact name match.

Brand	Named in answers	Rank 1 (of 20 generic answers)
Attention Insight	20 / 25	10
Neurons / Neurons Predict	12 of 20 generic (plus subject of Q3)	1
EyeQuant	11 of 20 generic	4 (all ChatGPT, plus table leads)
Brainsight	10 of 20 generic	2
VisualEyes	7 of 20 generic	1

Two runs were personalized by engine account memory and are tallied but flagged (R12, R13, R18, R22; detail held in engagement records).

E2. Question map

	ChatGPT	Claude	Perplexity	Gemini	Google AI Overviews
Q1 best attention heatmap tool	2nd	1st	1st	1st	1st
Q2 best predictive eye tracking software	2nd	1st	1st	1st	absent
Q3 Neurons Predict alternatives	2nd	1st	absent	absent	absent
Q4 cost / best service	2nd	absent	1st	2nd	1st
Q5 accuracy and reviews (about you)	subject	subject	subject	subject	subject

E3. What AI gets wrong

ID	Engine	What AI says	Reality	Traced source	Fix
ACC-02	Claude	Solo \$23/mo, 20 credits, 7-day trial; Team \$479/mo, 5 seats	Plans are Basic EUR 29 / Pro EUR 119 / Hero EUR 299; 14-day trial	Capterra listing , SUR-03 (confirmed)	FIX-01
ACC-03	Perplexity	"from about \$19/month"	Cheapest is EUR 23/mo billed annually; monthly Basic EUR 29 (\$31)	stale aggregator, SUR-08/SUR-12 family, untraced (blocked readers)	FIX-10
ACC-10	Gemini	"trained on over 70,000 real eye-tracking study results"	Client materials state ~30,800 study images and 5.5M+ fixations	no source carries it; drift from inconsistent figures	FIX-03
ACC-07	Claude	founded 2019	Sources conflict (2018 press vs 2019 Crunchbase); your confirmation pending	Crunchbase, SUR-06; press conflict SUR-07 (confirmed as source)	FIX-06

ACC-06/08/09	various	Extension install counts, review totals, Adobe free plugin	Platforms block automated readers; figures appear only as their listings show	listings (SUR-04, SUR-05, SUR-09)	FIX-07 , FIX-10
--------------	---------	--	---	-----------------------------------	---

The correct claims: current prices when read from your live site, the accuracy range as your own claim, watermark policy, results-in-seconds (ACC-01, ACC-04, ACC-05, ACC-11, ACC-12).

E4. How AI frames you

Engine	Descriptors attached to your name
ChatGPT	"best value", "easiest choice", "considerably cheaper", consistently second to EyeQuant for "professional" work
Claude	"the best all-around pick", "the default pick", "most widely adopted", "most established"
Perplexity	"most purpose-built and widely recommended", "go-to for fast design testing"
Gemini	"Top Overall for Designers", "best balance of accuracy, speed, and integration"
Google AI Overviews	"widely considered the best dedicated AI predictive heatmap tool", "Best Overall for Marketing & UX"

Net framing is strongly positive. The one recurring hedge is ChatGPT's value-tier framing: excellent but the budget option next to EyeQuant. The comparison and cost content ([FIX-04](#), [FIX-05](#)) is where that framing gets contested with checkable facts.

E5. Open territory

Gap	Question or gap	Why open	Fix
GAP-01	Neurons Predict alternatives	No comparison content on site; absent from cited roundups (CIT-12 to CIT-16)	FIX-04 , FIX-08
GAP-02	Best-software question on Google AI Overviews	GAIO cites competitor pages and a 2019 roundup (CIT-13, CIT-15)	FIX-05 , FIX-08
GAP-04	Category cost question	No cost guide; aggregator fragments fill the void	FIX-05 , FIX-10
GAP-05	One accuracy story	Four figure variants circulating plus one invention	FIX-03
GAP-09	Review recency	Trustpilot quiet ~12 months (SUR-05, CIT-08)	FIX-07

E6. Technical findings

Check	Status	Severity	Fix
robots.txt	All five AI crawlers allowed (SUR-13)	pass	-
Copy readable without JavaScript	100-150% of rendered copy present in raw HTML on every page (SUR-19)	pass	-
Pricing page identity	/pricing/ canonicals to /billing-plans/; /pricing/ absent from sitemap; both serve identical content (SUR-16, SUR-02)	high	FIX-02
Figma plugin identity	Three addresses for one topic, one 301 and two live (SUR-18)	medium	FIX-02
Structured data	Organization present; pricing page has none; no Product, Offer, or FAQ blocks anywhere (SUR-19)	medium	FIX-09
HTML tables	Zero real tables site-wide, including pricing (SUR-19, EX-05)	medium	FIX-09

llms.txt	404 (SUR-15)	low	FIX-11
Sitemap	Index with 10 child sitemaps, 57 pages listed; pricing missing, see above (SUR-16)	covered	FIX-02

E7. Records

Every report ID (Q, R, ACC, SUR, EX, CIT, GAP, FIX) resolves to a recorded evidence row held in MentionShare's engagement records, available on request.

Re-testing. AI answers move: cited sources shift month to month even when nothing changes on your side. Compare re-test numbers against the run-level detail, and treat single-run swings as noise. Fixes on live-retrieved surfaces (your pages, your profiles, articles that cite you) can show up in answers within days to weeks of a recrawl. Answers rooted in model training move on model-update timescales that nobody outside the AI labs controls.