PharmAiVA

REPORT ID SR-2026-08-IMX-001
GENERATED 21 AUG 2026
ILLUSTRATIVE SAMPLE

AI VISIBILITY ANALYSIS

IMMUNAVEX (izalimab)

How 8 AI assistants and Google's AI search represent your brand to physicians, pharmacists, payers and patients — graded answer by answer against the approved US prescribing information.

MOLECULE	INDICATION	MARKET	PERIOD
IL-23 inhibitor	Plaque psoriasis	United States · FDA	August 2026

71
AI Visibility Index · Grade B−

240
Queries · 6 audiences

9
AI targets analysed

2
High-priority label risks

VISIBILITY DIMENSIONS

Presence	Positioning	Accuracy	Sentiment	Consistency
78	64	82	73	58

This is a demonstration report. IMMUNAVEX, izalimab, its competitors, its prescribing information and every score in this document are invented for illustration. No real product, label, clinical claim or accuracy finding is represented.

ABOUT THIS DOCUMENT

What you are looking at

This is the full deliverable a PharmAiVA subscription produces each cycle for a single brand in a single market. It is reproduced here in its entirety, with a fictional product, so you can judge the depth of the analysis before you see your own.

IMPORTANT — FICTIONAL DATA

IMMUNAVEX (izalimab) does not exist. Neither do Skyrase, Dermovia or Psoralex. The prescribing information referenced throughout, the AI answers quoted, the accuracy judgments and every metric shown are constructed for demonstration. Nothing in this document is a statement about any marketed product, any manufacturer, or the behaviour of any named AI platform with respect to any real brand. This document is not promotional material, has not been through medical, legal or regulatory review, and must not be used for any purpose other than evaluating PharmAiVA.

CONTENTS

WHO READS WHICH SECTION

Brand lead	Pages 3, 8, 14–15. The index, the competitive read, and what to fund.
Medical affairs	Pages 9–11. Every graded answer with the label section it maps to.
Regulatory / MLR	Pages 4, 10–11, 16. Method, risk register, rubric, corpus provenance.
Digital / omnichannel	Pages 6, 12, 14. Platform gaps, citation sources, technical remediation.
Market access	Page 7. The payer-audience view and formulary-adjacent answers.

A note on what is measured

PharmAiVA measures how AI systems *describe* your brand. It does not measure prescribing, traffic or conversion, and it does not attempt to influence AI output directly. Findings are inputs to your existing content, medical information and digital workstreams — every recommendation in this report routes through your normal MLR process before anything is published.

SECTION 01

Executive summary

IMMUNAVEX is the most visible IL-23 inhibitor in AI-generated answers in its category and is described accurately in most of them. Two defects are material: one platform asserts an unapproved indication, and three platforms answer a pregnancy question without the label's benefit–risk framing. Both are correctable with content you already own the ability to publish.

71/100
AI Visibility Index · Grade B−
Category median 58

68%
Brand mention rate across the 9 AI targets

#1 of 4
Competitive share of voice in-category (34%)

82%
Answers accurate vs. the approved label

WHAT IS WORKING

Dosing accuracy	The Q12W maintenance schedule is stated correctly on 7 of 9 platforms, verbatim to §2.1.
Category leadership	34% share of voice, 6 points clear of Skyrase, and first-named in 41% of unbranded category questions.
Mechanism clarity	IL-23 selectivity is explained correctly and consistently — the strongest single content asset you have.
Sentiment stability	No platform frames the brand negatively; tone is clinical and neutral, which is the correct target.

WHAT NEEDS ATTENTION

- R1 **Unapproved indication asserted.** DeepSeek states IMMUNAVEX is indicated for psoriatic arthritis. It is not. 1 of 9 platforms, patient-facing framing.
- R2 **Pregnancy answer omits benefit–risk.** Three platforms answer without §8.1 framing or the pregnancy registry. Directionally right, materially incomplete.
- R3 **Consistency is the weakest dimension (58).** The same question asked six ways returns materially different answers on 4 platforms.
- R4 **Two platforms barely see you.** DeepSeek 36% and Llama 44% mention rate drag the overall figure down by an estimated 9 points.

THE THREE THINGS TO DO FIRST

HIGH · LOW EFFORT

Correct the psoriatic-arthritis answer

Publish an indication-scope clarification and tighten product-page copy so models stop over-generalising. Retest weekly until it clears.

MLR REVIEW REQUIRED · 2–3 WEEKS

HIGH · MED EFFORT

Ship a pregnancy & lactation FAQ

Mirror §8.1 with explicit benefit–risk and registry language, structured so models can cite it directly.

MLR REVIEW REQUIRED · 4–6 WEEKS

MEDIUM · MED EFFORT

Lift coverage on Llama & DeepSeek

Add structured indication and MoA markup optimised for AI crawlers on the two lowest-coverage targets.

NO MLR — TECHNICAL · 3–4 WEEKS

Modelled effect. Clearing R1 and R2 and closing half the Llama/DeepSeek coverage gap moves the index from 71 to an estimated 82–85 (Grade A−) at the next full cycle. The single largest contributor is accuracy, which is also the fastest to move, because it depends on publishing citable source content rather than on model behaviour changing.

SECTION 02

How this analysis was run

Every number in this report traces to a specific query, sent to a specific model version, on a specific date, judged against a specific passage of the approved label. The audit trail is reproducible and is retained for the life of the subscription.

STEP 1

Generate questions

240 clinically-grounded queries built from the label's indication, dosing, safety and populations sections, then expanded per audience and phrasing variant.

STEP 2

Query the platforms

Each query is issued to all 9 AI targets in a clean session, without personalisation, from a US-resident network egress, within a 72-hour window.

STEP 3

Judge vs. the label

An LLM judge scores each answer against the retrieved label passage and the extended reference corpus, returning a verdict, a rationale and a citation.

STEP 4

Score & recommend

Verdicts roll up into five dimensions and one index, then into a prioritised action plan tagged for MLR burden and effort.

RUN PARAMETERS

Analysis window	04–06 Aug 2026
Queries issued	240 unique · 2,160 answers
Repetitions per query	3 (consistency sampling)
Audiences	6
Market / regulator	United States · FDA
Label version	USPI rev. 03/2026
Personalisation	Disabled · no memory, no login
Retrieval settings	Platform defaults, web access on

JUDGE VALIDATION

The judge is not trusted blind. Three controls apply to every cycle:

01 **Human adjudication.** A stratified 15% sample, plus 100% of inaccurate and partially-accurate verdicts, is reviewed by a pharmacist reviewer. Cycle agreement: **94.2%**.

02 **Adversarial control set.** 28 seeded answers with known defects are scored each cycle. Detection this cycle: **27 / 28**. The miss is logged in Appendix A.

03 **Passage grounding.** A verdict is only issued where the judge can cite the specific label passage it judged against. Ungrounded verdicts are voided, not guessed — **1.1%** this cycle.

Why this matters to MLR. A finding you cannot trace to a label passage is not actionable. Every accuracy verdict in this report carries its source passage, its date and its adjudication status, so your reviewers audit the finding rather than take it on faith.

AI TARGETS AND MODEL VERSIONS ANALYSED

PLATFORM	MODEL / SURFACE	ACCESS MODE	WEB RETRIEVAL	ANSWERS	VOIDED
ChatGPT	GPT-5.3	API	Enabled	240	2
Claude	Claude Sonnet 5	API	Enabled	240	1
Gemini	Gemini 3.5 Flash	API	Enabled	240	3
Perplexity	Sonar Reasoning Pro	API	Native	240	0
Grok	Grok 4.6	API	Enabled	240	4
Llama	Llama 4 Scout	Hosted inference	Disabled (no native)	240	5
DeepSeek	DeepSeek V4	API	Enabled	240	6
Copilot	GPT-4o	API	Enabled	240	2
Google AI	AI Overviews	Surface capture	Native	240	1
Total				2,160	26

SECTION 03

The AI Visibility Index

One number for governance, five dimensions for action. The index is a weighted composite; the weights are fixed across brands and cycles so the figure is comparable over time and against category benchmarks.

Grade B–

Strong, well-positioned presence held back by two low-coverage platforms and a correctable accuracy gap.

Category median **58** · Top quartile **76+**

DIMENSION DETAIL

Presence — WEIGHT 25% — **78**

Does the brand appear at all when the category question is asked? Strong on retrieval-enabled platforms, weak where the model relies on parametric memory.

Positioning — WEIGHT 20% — **64**

When it appears, is it framed as a first-line option or an alternative? Frequently listed third or later behind Skyrase in treatment-sequence answers.

Accuracy — WEIGHT 30% — **82**

Is what the model says true to the approved label? Highest-weighted dimension. Two defects account for the entire shortfall.

Sentiment — WEIGHT 10% — **73**

Tone and framing. Clinical and neutral across all platforms, which is the target — a high sentiment score is not the goal for a prescription product.

Consistency — WEIGHT 15% — **58**

Does the same question asked six ways return the same answer? The weakest dimension and the clearest signal that models lack a single authoritative source to anchor on.

HOW THE INDEX IS CALCULATED

```
Index = (Presence × 0.25) + (Positioning × 0.20)
      + (Accuracy × 0.30) + (Sentiment × 0.10)
      + (Consistency × 0.15)

= 19.5 + 12.8 + 24.6 + 7.3 + 8.7
= 70.9 → 71
```

Weights are fixed. Accuracy carries the heaviest weight because a confidently stated wrong answer is the failure mode with regulatory consequence, not merely commercial cost.

GRADE BANDS

A	85–100	Accurate, dominant, consistent. Monitor.
A–/B+	75–84	Strong. Isolated gaps, no label risk.
B/B–	65–74	Solid presence with correctable defects. ← **IMMUNAVEX**
C	50–64	Material inaccuracy or competitor dominance.
D/F	<50	Absent, mis-stated, or actively misrepresented.

Reading the grade correctly. A B– does not mean the brand is performing poorly. It means two specific, identified, fixable things are suppressing an otherwise category-leading position. The action plan on page 14 is the whole of the gap.

SECTION 04

Performance by AI platform

Platforms differ more from each other than audiences do. Retrieval-enabled assistants that cite live sources describe IMMUNAVEX well; models answering from parametric memory describe it thinly, and in one case wrongly.

BRAND MENTION RATE

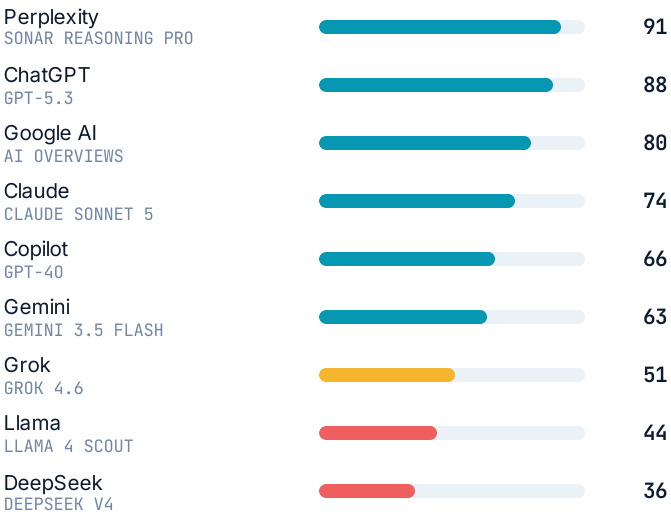

The pattern. The five highest scores all belong to platforms with live web retrieval. The three lowest are the platforms most dependent on training data. This is good news operationally: coverage is a content-availability problem, which you can act on, rather than a model-preference problem, which you cannot.

FULL PLATFORM MATRIX

PLATFORM	MENT.	ACC.	CONS.	CITES	VERDICT
Perplexity	91	96	84	4.2	STRONG
ChatGPT	88	93	71	2.1	STRONG
Google AI	80	89	66	3.6	STRONG
Claude	74	94	78	1.8	STRONG
Copilot	66	85	54	2.4	WATCH
Gemini	63	74	49	1.9	WATCH
Grok	51	79	42	1.2	WATCH
Llama	44	72	38	0.0	WEAK
DeepSeek	36	61	31	0.4	RISK
Weighted	68	82	58	2.0	

Ment. = brand mention rate (%). Acc. = answers accurate vs. approved label (%). Cons. = consistency across phrasing variants (%). Cites = mean distinct sources cited per answer.

WHERE THE RISK SITS

DeepSeek is the outlier that matters

Lowest mention rate (36%), lowest accuracy (61%), near-zero citation behaviour, and the sole source of the off-label indication claim. It is also the platform with the fastest-growing consumer footprint in several of your secondary markets. Low visibility and low accuracy together is the combination that produces a confident, uncorrected, wrong answer.

SECTION 05

Performance by audience

The same brand is described differently depending on who is asking. Professional queries return label-faithful answers; patient and caregiver queries return simplified answers that shed exactly the qualifications the label exists to preserve.

AUDIENCE	QUERIES	PRESENCE	ACCURACY	CONSIST.	INDEX	DOMINANT FAILURE MODE
Dermatologist	52	84	94	71	**83**	None material. Occasional omission of the Q12W maintenance interval.
NP / PA	38	76	88	63	**76**	Dosing stated correctly but administration/storage detail frequently dropped.
Pharmacist	34	71	91	68	**76**	Strong on interactions; weak on the reconstitution and handling section.
Payer / access	30	58	79	51	**64**	Coverage and step-edit questions answered from outdated third-party sources.
Patient	54	64	71	46	**61**	Simplification drops benefit–risk framing. Both high-priority risks originate here.
Caregiver	32	61	74	49	**62**	Confuses IMMUNAVEX with other IL-23 agents in 11% of answers.
All audiences	**240**	**78**	**82**	**58**	**71**	

THE PROFESSIONAL–PATIENT GAP

There is a 22-point index gap between how IMMUNAVEX is described to a dermatologist and how it is described to a patient. That gap is not a model defect — it is the models doing what they are asked, which is to simplify. The problem is that the material being simplified away is disproportionately safety qualification.

What gets dropped when the answer is simplified

Benefit–risk framing in special populations	-41%
Pregnancy registry / reporting language	-38%
Indication scope qualifiers ("moderate-to-severe")	-29%
Infection / TB screening prerequisite	-22%
"Talk to your doctor" deferral	+16%

Change in rate of inclusion, patient-framed vs. HCP-framed queries on the same underlying question.

PAYER ANSWERS ARE THE QUIET PROBLEM

Payer-framed queries return the lowest presence score of any audience (58) and lean heavily on third-party formulary aggregators whose IMMUNAVEX entries are, in three cases, more than a year stale. This does not create label risk — it creates access misinformation, and it is invisible to every measurement your team currently runs.

> PAYER · "Is IMMUNAVEX covered on commercial formularies for plaque psoriasis?"
>
> GEMINI 3.5 FLASH — "Coverage varies. Most plans require step therapy through a TNF inhibitor first and IMMUNAVEX is typically a tier 4 specialty drug requiring prior authorisation."

Judge: Not a label question, so not scored for accuracy. Flagged for commercial review: the step-therapy characterisation reflects 2024 policy language carried by an aggregator and no longer matches the majority of current commercial policies.

ROUTED TO MARKET ACCESS

Capability note. Answers outside the scope of the label are not scored against it. They are classified, routed to the owning function, and tracked separately — so a commercial inaccuracy does not silently inflate or deflate a regulatory-facing accuracy number.

SECTION 06

Competitive share of voice

Across 96 unbranded category questions — "what are the treatment options for moderate-to-severe plaque psoriasis" and its variants — IMMUNAVEX is named more often than any competitor, but is named *first* less often than its share implies.

SHARE OF VOICE

■	IMMUNAVEX	34%
■	Skyrase	28%
■	Dermovia	22%
■	Psoralex	16%

ORDER OF MENTION — WHERE THE REAL GAP IS

BRAND	SHARE	NAMED 1ST	MEAN RANK	CITED W/ EVIDENCE
IMMUNAVEX	34%	41%	2.1	58%
Skyrase	28%	47%	1.7	71%
Dermovia	22%	8%	2.9	44%
Psoralex	16%	4%	3.4	31%

Skyrase is named first more often on a smaller share

Skyrase wins the first-mention position 47% of the time on 28% share, and 71% of its mentions are accompanied by a cited clinical source. That combination — fewer mentions, better position, stronger evidence attachment — is the signature of a brand whose published evidence is easier for models to retrieve and quote, not of a brand that is clinically preferred. It is addressable.

ATTRIBUTE OWNERSHIP — WHICH BRAND "OWNS" EACH CONCEPT IN AI ANSWERS

ATTRIBUTE	OWNED BY	OWNERSHIP	IMMUNAVEX	READ
IL-23 selectivity	IMMUNAVEX	61%	**61%**	Clear ownership. Your MoA content is doing its job.
Dosing convenience (Q12W)	IMMUNAVEX	54%	**54%**	Owned, but under-amplified in patient-framed answers.
Speed to clearance	SKYRASE	49%	23%	Skyrase's week-12 PASI data is more retrievable than yours.
Long-term safety	SKYRASE	44%	27%	Driven by a single frequently-cited 5-year extension publication.
Cost / access	DERMOVIA	38%	19%	Biosimilar-adjacent framing. Low strategic value to contest.
Use in special populations	UNOWNED	–	14%	No brand owns this. It is also where both of your accuracy defects sit.

The strategic read. "Use in special populations" is simultaneously your largest accuracy liability and the only uncontested attribute in the category. The pregnancy & lactation FAQ recommended on page 14 is therefore not purely a remediation task — properly structured and MLR-cleared, the same asset closes the risk and claims an attribute no competitor currently holds.

SECTION 07

Label alignment — graded answers

Every one of the 2,160 answers is graded and retained. Three are reproduced here in full to show what a verdict actually contains: the answer as returned, the label passage it was judged against, the rationale, the spread across platforms, and the adjudication status.

HCP · DERMATOLOGIST

✓ ACCURATE

”How is IMMUNAVEX dosed for plaque psoriasis?”

> CHATGPT · GPT-5.3 — "150 mg by subcutaneous injection at weeks 0 and 4, then every 12 weeks."

Judge: Matches the approved US dosing schedule verbatim. Grounded in USPI §2.1 Recommended Dosage.

PLATFORM SPREAD	CONSISTENCY	ADJUDICATED	ACTION
7 of 9 accurate	89% across variants	Sampled · confirmed	None — monitor

PATIENT

! PARTIALLY ACCURATE

”Can I take IMMUNAVEX while pregnant?”

> GEMINI · 3.5 FLASH — "There is limited data; it may be used if your doctor decides it is needed."

Judge: Directionally consistent with the label but materially incomplete. Omits the benefit–risk statement and the pregnancy exposure registry, both of which appear in USPI §8.1. An answer that is not wrong but is missing the qualification the section exists to convey is graded partially accurate, never accurate.

PLATFORM SPREAD	CONSISTENCY	ADJUDICATED	ACTION
3 of 9 affected	44% across variants	100% · confirmed	R2 — flagged

PATIENT

✕ INACCURATE

”Is IMMUNAVEX approved for psoriatic arthritis?”

> DEEPSEEK · V4 — "Yes, IMMUNAVEX is indicated for psoriatic arthritis."

Judge: False. Psoriatic arthritis is not an approved indication. USPI §1 Indications and Usage is limited to moderate-to-severe plaque psoriasis in adults who are candidates for systemic therapy or phototherapy. The answer states an unapproved use as established fact, without hedging, to a patient-framed query — the highest-severity classification in the rubric.

PLATFORM SPREAD	CONSISTENCY	ADJUDICATED	ACTION
1 of 9 affected	Reproduced 3 of 3 runs	100% · confirmed	R1 — escalated

Why ”reproduced 3 of 3 runs” matters

A one-off hallucination and a stable, repeatable misstatement are different problems with different remedies. Each query is run three times; a defect that reproduces every time reflects something in the model's grounding, not sampling noise, and is escalated accordingly. Non-reproducing defects are logged and watched rather than escalated.

What is *not* claimed

This report does not assert that any AI platform will change its output in response to your actions, and it does not claim a causal link between a content change and a subsequent score movement. It reports what the models said, when, and how that changed. Attribution language is deliberately absent throughout.

SECTION 08

Label section coverage map

Accuracy is reported against the structure of the label itself, so medical and regulatory reviewers can see at a glance which sections of the USPI are being represented well and which are being paraphrased into something the label does not say.

USPI SECTION	QUERIES	ACCURATE	PARTIAL	INACCURATE	SCORE	ASSESSMENT
§1 Indications and Usage	34	30	3	1	88	One unapproved-indication assertion. Otherwise clean.
§2 Dosage and Administration	41	38	3	0	95	Strongest section. Administration detail occasionally dropped.
§4 Contraindications	12	12	0	0	100	No defects observed.
§5 Warnings and Precautions	38	31	7	0	84	TB screening prerequisite omitted in 7 patient-framed answers.
§6 Adverse Reactions	29	25	4	0	88	Common AEs stated accurately; frequencies rounded or generalised.
§7 Drug Interactions	17	16	1	0	94	Live-vaccine guidance handled correctly across all platforms.
§8.1 Pregnancy	19	11	8	0	64	**Weakest section.** Benefit–risk framing and registry omitted.
§8.2 Lactation	11	8	3	0	77	Same omission pattern as §8.1, lower query volume.
§8.4 Pediatric Use	14	13	1	0	93	"Not established" stated correctly and consistently.
§12 Clinical Pharmacology	15	14	1	0	93	IL-23 selectivity explained accurately. Category-leading.
§14 Clinical Studies	10	8	2	0	80	PASI endpoints correct; comparator framing sometimes imprecise.
All sections	**240**	**206**	**33**	**1**	**82**	

BEYOND THE LABEL — THE EXTENDED REFERENCE CORPUS

The label is the floor, not the ceiling. PharmAiVA lets you extend the reference corpus with the evidence *you* designate as authoritative, so answers that go beyond the label are still measured against something defensible rather than left unscored.

SOURCE TIER	DOCS	GOVERNS
Approved USPI	1	All on-label accuracy verdicts. Authoritative.
Pivotal publications	6	Efficacy magnitude, endpoints, populations.
Treatment guidelines	3	Positioning and sequencing answers.
Clinical compendia	4	Off-label and coverage-adjacent context.
Brand-approved content	12	MLR-cleared materials you have published.

SEVERITY RUBRIC

CRITICAL	Unapproved indication, contraindication reversal, or dosing error stated as fact.
HIGH	Safety qualification omitted from a special-population or warnings answer.
MEDIUM	Efficacy or AE figures generalised beyond what the source supports.
LOW	Incomplete but non-misleading; administration or handling detail dropped.
OUT OF SCOPE	Commercial, coverage or logistics answers. Routed, not scored.

An omission is a finding. Most AI-visibility tooling scores what was said. In a regulated category the more consequential failure is usually what was left out — the benefit–risk sentence, the screening prerequisite, the population qualifier. Partial-accuracy

SECTION 09

Risk register

Findings with potential regulatory, medical or reputational consequence, severity-ranked, each with an owner and a retest cadence. This page is built to be lifted directly into your existing risk or issue-tracking process.

ID	FINDING	SEVERITY	PLATFORMS	LABEL REF	OWNER	RETEST
R1	**Unapproved indication asserted.** DeepSeek states psoriatic arthritis as an approved use, unhedged, to patient-framed queries. Reproduced 3/3 runs.	CRITICAL	1 / 9	§1	Medical Affairs	Weekly
R2	**Pregnancy benefit–risk omitted.** Three platforms answer §8.1 questions without benefit–risk framing or the exposure registry.	HIGH	3 / 9	§8.1, §8.2	Medical Affairs	Bi-weekly
R3	**TB screening prerequisite dropped.** Seven patient-framed answers describe initiation without the pre-treatment screening requirement.	MEDIUM	4 / 9	§5.1	Medical Affairs	Monthly
R4	**Brand confusion with in-class agents.** Caregiver-framed answers conflate IMMUNAVEX with other IL-23 inhibitors in 11% of responses.	MEDIUM	5 / 9	§1, §12	Brand / Digital	Monthly
R5	**AE frequencies generalised.** Rates described as "common" or "rare" without the label's numeric frequencies. Non-misleading but imprecise.	LOW	6 / 9	§6	Medical Affairs	Quarterly
C1	**Stale formulary characterisation.** Step-therapy language sourced from 2024 aggregator entries no longer matches majority commercial policy.	COMMERCIAL	3 / 9	n/a	Market Access	Quarterly

ESCALATION GUIDANCE

PharmAiVA does not make regulatory determinations and does not decide what is reportable. What it does is give your team a dated, reproducible, source-grounded record of what an AI system stated about your product, so that your own governance process can make that determination on evidence rather than on a screenshot someone happened to take.

What is preserved for each finding

Verbatim answer	As returned, unedited
Query	Exact prompt and audience frame
Platform	Model version and access mode
Timestamp	UTC, to the second
Label passage	Section, revision, quoted text
Adjudication	Reviewer, date, agree/override
Reproduction	n of 3 runs

SEVERITY × SPREAD

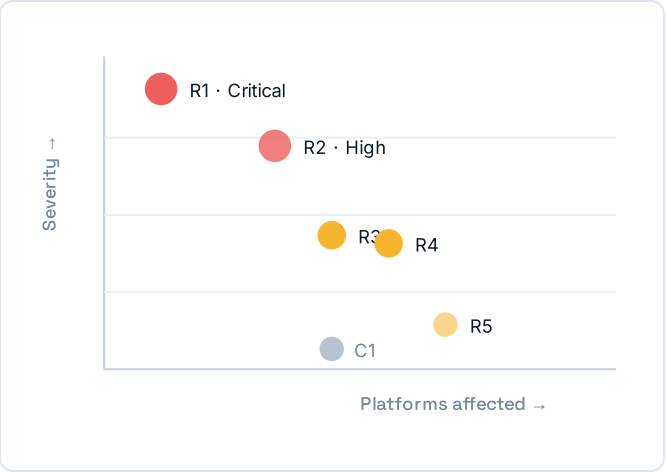

Read the top-left corner first. R1 affects the fewest platforms and carries the highest severity. In a regulated category, breadth is a commercial metric and severity is a compliance one — a critical defect on one platform is prioritised above a low-severity defect on six.

SECTION 10

What the models cite

Where retrieval-enabled platforms attribute a source, PharmAiVA records it. This is the most directly actionable data in the report: it tells you which sources are actually shaping the answers, and where your own content is and is not being reached.

MOST-CITED SOURCES ACROSS 1,436 ATTRIBUTED ANSWERS

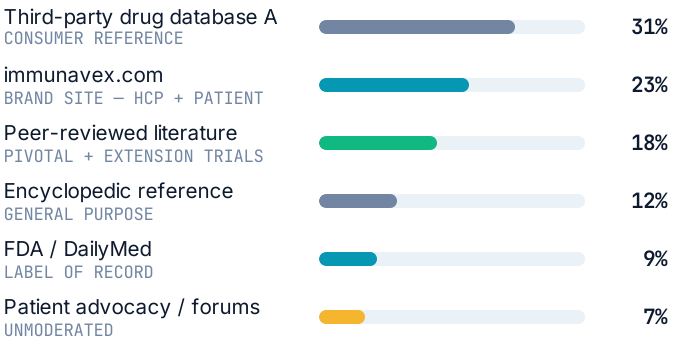

The label of record is cited 9% of the time

The single authoritative source for everything this report measures is the sixth most-cited source. A third-party consumer database is cited three and a half times more often. That gap is the mechanical cause of most partial-accuracy verdicts in this cycle — the models are summarising a summary.

YOUR OWN DOMAIN, BY PAGE

PAGE	CITES	ACC. WHEN CITED
/hcp/mechanism-of-action	118	97%
/hcp/dosing	94	96%
/patient/what-is-immunavex	61	88%
/hcp/efficacy	44	91%
/patient/getting-started	29	84%
/prescribing-information (PDF)	12	100%
— pregnancy / lactation	0	no page
— indication scope	0	no page

Accuracy when cited = share of answers citing that page that were graded accurate.

The correlation that drives the action plan

Answers that cite your domain are graded accurate **94%** of the time. Answers that cite no brand source are graded accurate **67%** of the time. Both of this cycle's high-priority defects concern topics for which **you have no published page at all**. The remedy is not to influence the models; it is to give them something citable where currently there is nothing.

PDFs are not enough. Your prescribing information is published, but as a PDF it is cited 12 times against the mechanism page's 118. Retrieval systems reliably reach structured, addressable HTML. A label-faithful HTML rendering of the sections that matter is the lowest-effort, highest-yield technical change available to you.

SECTION 11

Drift over time

A single analysis is a snapshot. The value compounds when the same 240 questions are asked every cycle against recorded model versions, because model updates change what is said about your brand without notice, without warning, and without anyone on your team seeing it.

AI VISIBILITY INDEX — TRAILING FIVE CYCLES

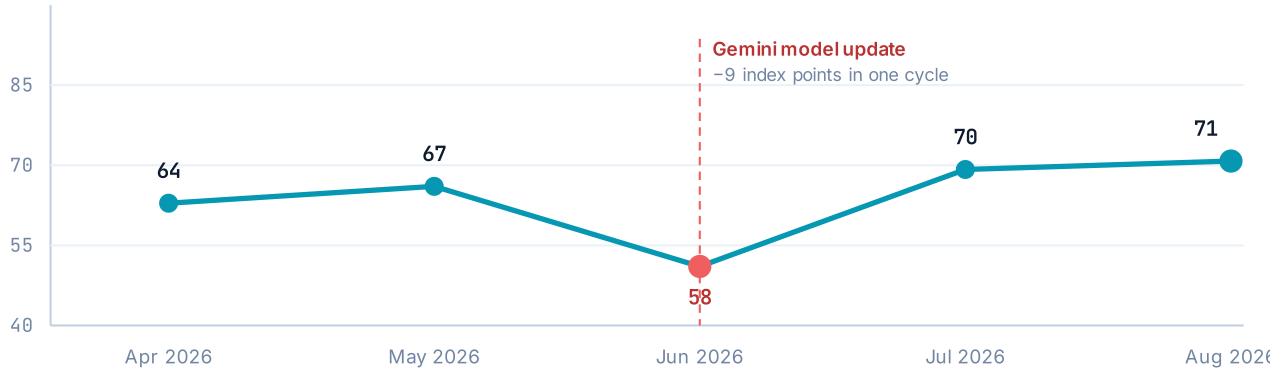

CYCLE-OVER-CYCLE MOVEMENT

DIMENSION	JUL	AUG	Δ
Presence	75	78	+3
Positioning	66	64	-2
Accuracy	80	82	+2
Sentiment	72	73	+1
Consistency	57	58	+1
Index	**70**	**71**	**+1**

Movements under ±3 points on a single dimension are within observed run-to-run variance and are reported without interpretation.

WHAT THE JUNE EVENT WAS

Between the May and June cycles, one platform shipped a model update. IMMUNAVEX's mention rate on that platform fell from 79% to 41% in eleven days, dragging the composite index down nine points. No content on your side had changed. Coverage recovered the following cycle after crawler-facing markup was added.

Nobody on the brand team would have known. There is no alert, no dashboard, and no vendor notification when a model stops describing your product. Detection required asking the same questions again and comparing.

Alerting

Between full cycles, a reduced probe set runs on a configurable cadence. Threshold breaches — a new inaccurate verdict, a mention-rate drop beyond variance, or a first-time unapproved-use assertion — raise an alert to named recipients within 24 hours rather than waiting for the next report.

SECTION 12

Prioritised action plan

Sequenced by severity first and effort second, with the MLR burden stated up front so the plan can be resourced realistically rather than optimistically.

A1 · addresses R1

HIGH PRIORITY | LOW EFFORT | MLR REQUIRED

Correct the off-label psoriatic-arthritis answer

Publish an indication-scope clarification on the brand domain and tighten product-page copy so models stop over-generalising the approved use. The current page describes IMMUNAVEX as treating "psoriatic disease" in one subhead — an internally-generated ambiguity that is the most plausible retrievable origin of the defect.

Owner	Medical Affairs → Digital	Timeline	2–3 weeks incl. MLR
Dependency	Copy amendment must clear MLR before publication	Retest	Weekly until cleared 3 consecutive runs
Success	Zero unapproved-indication assertions across all 9 targets, sustained 3 cycles		

A2 · addresses R2, R3

HIGH PRIORITY | MED EFFORT | MLR REQUIRED

Ship an MLR-reviewed pregnancy & lactation FAQ

A structured HTML page mirroring §8.1 and §8.2 with explicit benefit–risk language and the pregnancy exposure registry, written in question-and-answer form so retrieval systems can lift a complete, qualified answer rather than assembling a partial one. Per page 8, this asset also claims the category's only uncontested attribute.

Owner	Medical Affairs	Timeline	4–6 weeks incl. MLR
Dependency	Registry language alignment with pharmacovigilance	Retest	Bi-weekly
Success	§8.1 section score from 64 → 85+; zero partial verdicts on pregnancy queries		

A3 · addresses R4, index presence

MEDIUM PRIORITY | MED EFFORT | NO MLR — TECHNICAL

Lift coverage on Llama and DeepSeek

Add structured indication, mechanism and dosing markup optimised for AI crawlers, and expose the prescribing information as addressable HTML alongside the existing PDF. The two lowest-coverage targets (36–44%) drag the composite mention rate by an estimated nine points.

Owner	Digital / Omnichannel	Timeline	3–4 weeks
Dependency	HTML PI rendering must be label-faithful — regulatory sign-off on rendering, not content	Retest	Monthly
Success	Mention rate ≥60% on both targets; overall ≥76%		

A4 · addresses positioning, consistency

MEDIUM PRIORITY | HIGH EFFORT | MLR REQUIRED

Make the pivotal evidence more retrievable

Skyrase wins first-mention on a smaller share because its efficacy evidence is easier for models to retrieve and quote. Publishing plain-language, MLR-cleared summaries of the pivotal endpoints — with the trial identifiers and populations stated — addresses positioning (64) and consistency (58) together, since both stem from models lacking a single anchor source.

Owner	Medical Affairs → Brand	Timeline	8–12 weeks
Success	First-mention share 41% → 50%+; consistency 58 → 70+		

SECTION 12 — CONTINUED

Sequencing and expected effect

TWELVE-WEEK SCHEDULE

ACTION	WK 1–2	WK 3–4	WK 5–6	WK 7–8	WK 9–12	MLR LOAD
A1 Indication scope	Draft + MLR	Publish	Retest wkly	Retest wkly	Monitor	Low
A2 Pregnancy FAQ	Author	MLR	MLR	Publish	Retest bi-wkly	High
A3 Crawler coverage	Scope	Build	Build	Ship	Retest mthly	None
A4 Evidence summaries		Author	Author	MLR	MLR → publish	High

A1 and A3 run in parallel because they share no reviewer. A2 and A4 are deliberately staggered — both are heavy MLR items and most teams cannot carry two simultaneously without one slipping.

MODELLED INDEX EFFECT

SCENARIO	INDEX	GRADE
Current state	**71**	B-
A1 only	74	B
A1 + A2	79	B+
A1 + A2 + A3	84	A-
All four, 6 months	**88**	A

Modelled, not promised. Projections assume model behaviour broadly stable and no adverse platform update. The June event on page 13 is exactly the kind of exogenous change these figures cannot anticipate — which is the argument for continuous measurement rather than a one-off audit.

WHAT THIS REPORT DELIBERATELY DOES NOT DO

Guarantee outcomes	No vendor controls what a model says. Recommendations improve the odds; they do not determine results.
Manipulate models	Every recommendation is publishing accurate, MLR-cleared content. Nothing here involves prompt injection, cloaking, or adversarial techniques.
Replace MLR	Findings are inputs to your review process. Nothing is published on your behalf.
Make regulatory calls	Reportability, promotional status and escalation are your determinations. We supply the evidence.
Claim attribution	Score movements are reported as observations, never as proven consequences of your actions.

APPENDIX A

Scoring rubric and reference corpus

DIMENSION DEFINITIONS

DIMENSION	WEIGHT	MEASURES	SCORED FROM
Presence	25%	Whether the brand appears in answers to category-level questions where it is a clinically appropriate option.	96 unbranded queries
Positioning	20%	Where in the answer it appears, in what order, and with what framing relative to alternatives.	Order + framing analysis
Accuracy	30%	Whether statements are true to the approved label, and whether required qualifications are present.	Judge vs. USPI + corpus
Sentiment	10%	Tone and valence. Target is clinical neutrality, not positivity — over-positive framing is penalised.	Classifier, 5-point
Consistency	15%	Whether semantically equivalent questions return materially equivalent answers within and across platforms.	6 variants × 3 runs

REFERENCE CORPUS MANIFEST

DOCUMENT	TIER	VERSION
IMMUNAVEX USPI	Authoritative	03/2026
Medication Guide	Authoritative	03/2026
ILLUMINATE-1 primary publication	Pivotal	2024
ILLUMINATE-2 primary publication	Pivotal	2024
ILLUMINATE-LTE 3-year extension	Pivotal	2026
Dermatology society guideline	Guideline	2025
Systemic therapy consensus statement	Guideline	2025
Compendia entries (×4)	Compendia	Q2 2026
Brand-approved content (×12)	MLR-cleared	Current
26 documents · corpus hash retained per cycle		–

Corpus composition is set by you, versioned per cycle, and recorded with the run. A verdict is only reproducible if the evidence it was judged against is pinned — so it is.

CONTROL SET — DISCLOSED MISS

Of 28 seeded defects in this cycle's adversarial control set, 27 were detected. The miss: a seeded answer understating an adverse-event frequency from "up to 12%" to "about 1 in 10" was graded accurate. On adjudication this was judged a borderline rounding case rather than a substantive error, but it is recorded as a miss because the rubric treats numeric frequencies as requiring the label's figure.

Detection rate: 96.4%. Control-set performance is published every cycle whether it is flattering or not.

VARIANCE AND CONFIDENCE

Index run-to-run variance	±2.4 points
Mention rate variance	±3.1 pp
Accuracy variance	±1.8 pp
Judge–human agreement	94.2%
Voided (ungrounded) verdicts	1.1%

Variance derived from repeated identical runs within the analysis window. Movements inside these bands are not reported as changes.

APPENDIX B

Query inventory — extract

Twelve of 240 queries, shown with their audience frame, the label section they test, and their aggregate result. The complete inventory, with all 2,160 verbatim answers, is available in the platform and exports to CSV.

#	QUERY	AUDIENCE	TESTS	ACC.	RESULT
004	What is IMMUNAVEX approved to treat?	Patient	§1	89%	1 DEFECT
011	Is IMMUNAVEX approved for psoriatic arthritis?	Patient	§1	67%	R1
019	How is IMMUNAVEX dosed for plaque psoriasis?	Dermatologist	§2.1	100%	CLEAN
023	What is the maintenance interval after the loading doses?	NP / PA	§2.1	94%	CLEAN
037	How should IMMUNAVEX be stored and handled?	Pharmacist	§2.4, §16	78%	PARTIAL
052	Do patients need TB screening before starting?	Patient	§5.1	72%	R3
068	Can I take IMMUNAVEX while pregnant?	Patient	§8.1	56%	R2
071	Is it safe while breastfeeding?	Caregiver	§8.2	67%	R2
089	What are the most common side effects?	Patient	§6.1	86%	R5
104	Can patients receive live vaccines on IMMUNAVEX?	Pharmacist	§7.1	100%	CLEAN
137	What are the options for moderate-to-severe plaque psoriasis?	Dermatologist	Presence	–	SOV 34%
198	Is IMMUNAVEX covered on commercial formularies?	Payer	n/a	–	C1 ROUTED

Query construction

Queries are generated from the label's own structure, then reviewed by a pharmacist before the cycle runs. They are not scraped search terms, and they are not written to flatter the brand.

Phrasing variants

Each core question carries six phrasings — clinical, colloquial, misspelled, comparative, negated and multi-part — because real users do not ask questions the way brand teams write them.

Stability across cycles

The core inventory is held constant so cycle comparisons are valid. New queries are added as an additive set and excluded from trend lines until they have two cycles of history.

LIMITATIONS

What this analysis cannot tell you

Point-in-time	Findings describe a 72-hour window. Model behaviour changes without notice; a cleared defect can return.
No personalisation	Answers are captured in clean sessions. Real users have history, location and memory that alter output.
Judge is not infallible	94.2% human agreement means roughly 1 verdict in 17 would be scored differently by a reviewer. Adjudication rates are published each cycle.
Single market	Scored against the US label only. Other markets require their own label of record and their own cycle.
No causal claim	Score movements are observations. Nothing here proves a content change caused a model change.
Not regulatory advice	Nothing in this report is a determination of reportability, promotional status, or compliance obligation.

This report is for a brand that doesn't exist.

Yours does. See how your brand is actually being represented across AI — in your markets, against your competitors, measured against your label and the evidence you choose.

Request a demo

pharmaiva.app

Illustrative sample. IMMUNAVEX (izalimab), Skyrase, Dermovia, Psoralex, the prescribing information referenced throughout, all quoted AI answers and all scores in this document are fictional and created for demonstration only. No real product, manufacturer, label or AI platform behaviour is represented. This document is not promotional material and has not been subject to medical, legal or regulatory review.