

Two search markets, one country

We asked Google the same Saudi healthcare questions in Arabic and in English, then resolved every source it cited. The two languages are not two versions of one result. They are two different markets, and most providers are only competing in one of them.

VOCTOS research. Panel captured 9 and 10 September 2026. Method, raw data and limitations published in full.

0.145

Average overlap between the sources behind the Arabic answer and the sources behind the English answer, across **34 matched question pairs**. A score of 1.0 would mean identical sources. **12 pairs shared nothing at all.**

75	628	250	3.8%
captures	citations resolved	distinct domains	unresolved

What we actually did

This is a measurement, not an opinion piece. We asked matched pairs of real questions, one in Gulf Arabic and one in English, on the same day, from the same Riyadh location, logged out. For every answer we extracted each cited link and followed it to the page it actually opens, because the engine hides

destinations behind redirect tokens. 628 citations were resolved that way. 25 could not be, and they stay in the count rather than being quietly dropped.

The panel is Saudi healthcare: hospitals, clinics, laboratories, test prices, symptoms. The findings below are about that market. We say so rather than implying they generalise to yours, and the same protocol can be run on any sector.


Finding one: the two languages answer from different internets

Take one question. Best hospital in Jeddah, asked both ways within minutes. This is a high agreement pair by the standards of this dataset, and it still turns over half its sources.

Arabic prompt	5	English prompt	4
1	imc.med.sa	1	imc.med.sa
2	dsfhjeddah.fakeeh.care	2	dsfhjeddah.fakeeh.care
3	kch.sa	3	kch.sa
4	alsalamahospital.com	4	dsah.sa
5	kfshrc.edu.sa		

Shared domains: 3. Overlap 0.500. Resolved by hand from the capture log, one row per destination.

Across the whole panel the average is 0.145. Ranking well in English tells you very little about whether you appear in Arabic. They are separate battles fought against separate opponents.




Domain overlap for every matched pair, sorted

Finding two: in Arabic, the institutional wall is missing

In English, 27.7% of cited sources are medical institutions, publishers or government bodies. Mayo Clinic, the NHS, Cleveland Clinic. Those are almost impossible for a private provider to outrank.

In Arabic that share collapses to 10.7%, a gap of 17.0 points, while commercially owned pages carry 67.6% of the set. The Arabic answer is mostly built out of company websites. That is not a warning. It is the opening.



Who owns the cited page, Arabic against English

Finding three: ask "who is the best" in Arabic and institutions vanish

We isolated the commercial questions, the ones a patient with a wallet actually types. Best hospital, best doctor, best clinic, best lab.

74.7% of the Arabic sources on those questions are companies and 14.7% are booking directories. Institutions and government together account for 1.3%, which is 1 citation of 75. On the Arabic best-doctor question the engine drew nine citations from exactly two domains, both appointment platforms. Whoever those platforms list is who gets recommended.

Asked in Arabic

1.3%

1 / 75

Asked in English

7.6%

5 / 66

Institutional share on best-X questions

Finding four: for "best hospital" the answer is built from the map, not from articles

On the hospital questions almost every destination we resolved came from the places strip beside the answer rather than from a citation inside the text. The engine was reading business entities, not blog posts. Which means the lever on those questions is your entity data, your profile, your categories and your reviews, and not another article. Two of the thirty three sources on those four questions were inline citations. The rest were business listings.

Finding five: being cited once is not a position

We re asked the same questions 24 hours later, changing nothing. Mean agreement with the previous day was 0.602, and one pair held only 0.231. On one question the engine dropped the national accreditation body it had cited the day before and picked a different government portal instead. Anyone who shows you a single screenshot of an AI answer is showing you one day.

What this means if you are a provider in this market

1 Treat Arabic as a separate channel with its own inventory.

Only 45 of 250 domains in this dataset were cited in both languages. Translating your English pages is not a strategy, because the Arabic answer is drawing on a different set of pages entirely.

2 On commercial questions, fix your entity before your blog.

When the answer is assembled from the places strip, your profile data, categories, hours, reviews and the consistency of your name across directories are doing the ranking. Content cannot reach that surface.

3 Get onto the platforms the engine already trusts.

Booking directories carry 14.7% of Arabic commercial citations. Being listed, accurately and completely, is often faster than earning the citation yourself.

4 Measure repeatedly or do not measure at all.

Given the day to day movement we recorded, a single check tells you almost nothing. Position here is a distribution, not a number.

Want to know where you actually stand?

We will run your brand and your three closest competitors through the same protocol, in both languages, and show you every question where an AI assistant recommends someone else. You get the question list, the resolved sources, and the specific reason each competitor was chosen over you.

No rankings are promised. What is promised is that you will see exactly what the engine sees.

[Request an AI visibility audit](#)

Limitations, stated plainly. This panel is Saudi healthcare and is weighted toward Riyadh and Jeddah. Most captures are Google AI Mode; a smaller ChatGPT arm is reported separately. Coding was done by one analyst. Delta Medical Laboratories is a VOCTOS client and appears in the data, which is disclosed rather than removed. Full method, raw CSV and every limitation are published with the study.