ELNIQ GEO | queseo.es

RESEARCH INTO AI VISIBILITY

One question, many lists

How recommendations of SEO consultants in Valencia change when the same question is repeated to ChatGPT

200	198	26
ATTEMPTS	VALID LISTS	CANDIDATES

THE SAME QUESTION IN TWO ENVIRONMENTS

100 attempts in free ChatGPT in the browser, without signing in.
100 attempts through the OpenAI API, with web search enabled.

A STUDY BY ELNIQ AND QUESEO

Nina S. Dzhezher
ELNIQ SL (https://www.elniq.com/)

Eric Jorge Seguí Parejo
Queseo (https://queseo.es/)

Data collected on 8 and 9 November 2026
11 November 2026

QUICK READ

What this test shows

A single answer gives a partial view of the professionals ChatGPT may recommend. Repeating the same question changes the names and their order. The distribution of visibility also differs between the browser and the API.

Finding	Observed result
The ordered list almost never repeats	94 different lists in the browser and 95 in the API.
The environments show different options	A list from each environment shares an average of 1.77 of its five names.
Appearing and coming first are different measures	Queseo is the most frequent brand in the API; Jesús Porres tops more lists.
Sources and listings accompany different distributions	There are clear associations, but this test does not isolate their causes.
Content is also part of visibility	Specialties, links, reviews and possible identity mismatches should be checked.

THE SAMPLE ANALYSED

There were 200 recorded attempts with the same question: 100 in free ChatGPT in the browser, without signing in, and 100 through the OpenAI API with web search. Of these, 198 produced a list of five candidates. Counts are calculated directly from the two original Excel files from the test.

How to read it. The main text explains the test and its implications. The complete tables cover all 26 candidates. The appendices detail the stability calculations and uncertainty of the estimates.

01 / WHAT WAS DONE

The same question in two environments

In the browser, questions are asked on the ChatGPT website. The API lets a program send a question to OpenAI and receive an answer. Both access routes were tested. The Spanish prompt is translated below.

> **«Recommend 5 SEO consultants in Valencia, Spain. Briefly explain why you recommend each one.»**

Condition	Browser	API
Access	Free ChatGPT, without signing in. New chat for each attempt.	Requests from a program, with web search enabled.
Attempts and valid lists	100 attempts; 98 lists.	100 attempts; 100 lists.
Model	Not identified.	gpt-5.6-luna, the identifier requested and returned in the log.
Mainland Spain time	8 Nov, 21:34; until 9 Nov, 00:35.	8 Nov, 21:32:59 to 22:03:42.
Information retained	Text, site and page logs, listings and response elements.	Text, links, usage and number of search calls.

Browser attempts 28 and 44 stopped at the search message. They are excluded from the recommendation counts. List 85 retains visible text but lacks raw text and complete search logs.

What counts as a candidate

We count the five numbered names in each main list. Variants identifying the same person or brand are grouped. For example, **«Eric Jorge Seguí Parejo» and «Queseo» count as one candidate**. This gives 490 positions in the browser and 500 in the API. Candidates include professionals and teams.

GEO means Generative Engine Optimization: work on visibility in generative engine responses. This study observes that visibility; no intervention was applied to increase it. [3]

02 / REPEATING THE QUESTION

Names and order change

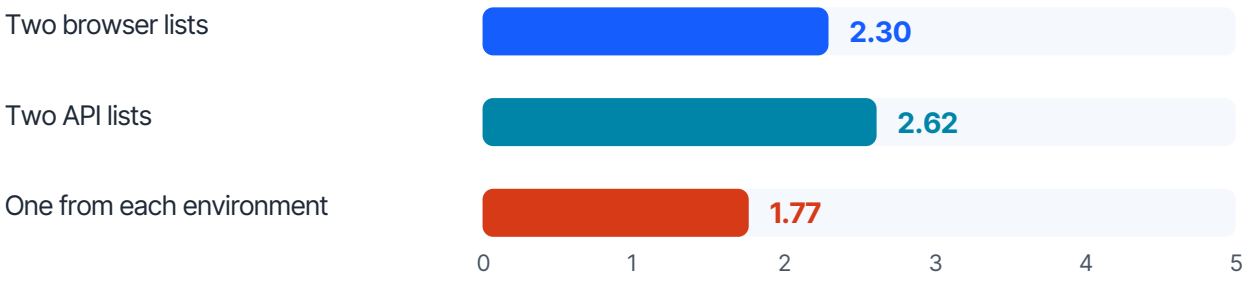
There are two ways to compare lists. A **combination** contains the same five names, even if their order changes. An **ordered list** also requires each name to occupy the same position.

A SIMPLE EXAMPLE

Ana, Bruno, Carla, Diego and Elena form a combination. If Bruno moves to first place and Ana to second, the combination stays the same, but the ordered list changes. If another person joins, the combination also changes. This is a fictional example.

What was counted	Browser 98 lists	API 100 lists
Different candidates	20	23
Combinations of five, ignoring order	68	67
Different lists when order is included	94	95
Answers whose ordered list does not appear again	91	90

How similar two answers are

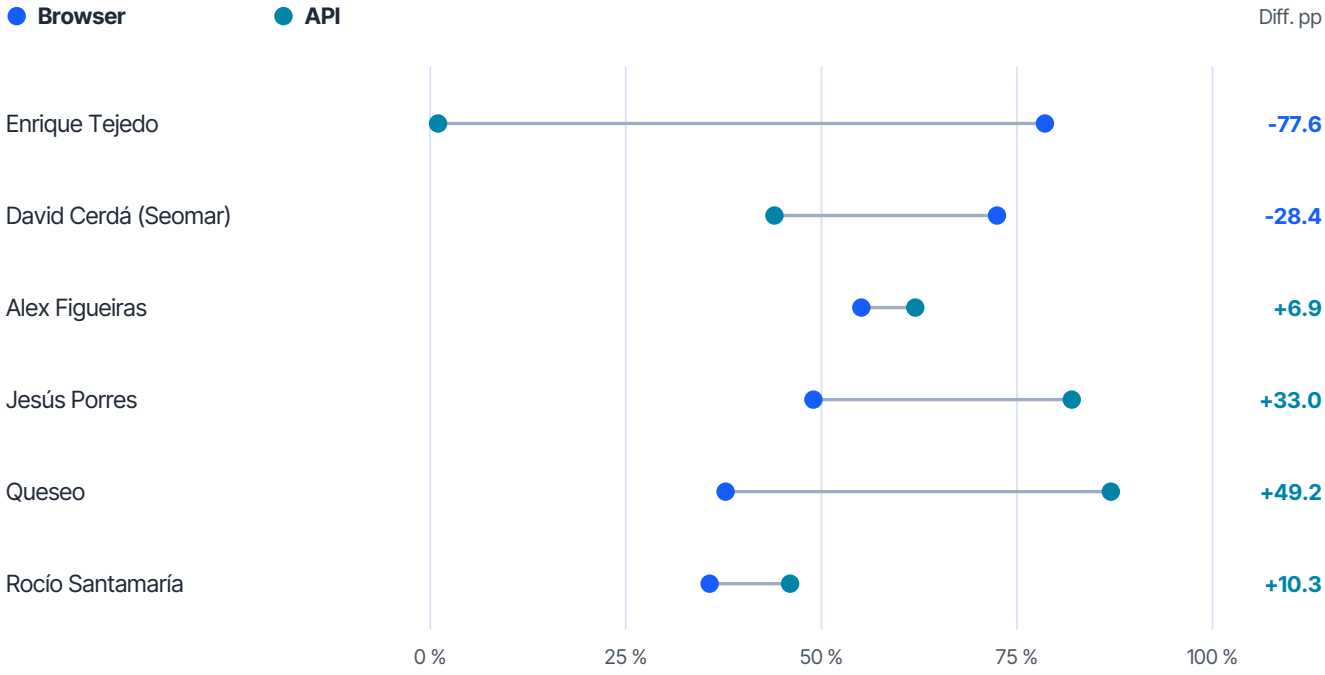


The bars count shared names, out of a maximum of five. All possible pairs were compared: 4,753 within the browser, 4,950 within the API and 9,800 across environments. No cross-environment list was identical in both names and order.

03 / WHO RECEIVES VISIBILITY

The distribution changes across environments

Presence is the percentage of lists in which a candidate appears in any of the five positions. Percentages are compared because there are 98 valid browser lists and 100 API lists.



Each point shows the brand's presence in one environment. The number at the right is the API minus browser difference, in percentage points. The five most frequent candidates in each environment are shown, six when combined.

How to read a difference

Queseo appeared in **37 of 98** browser lists (37.8%) and **87 of 100** API lists (87.0%): a difference of 49.2 points. Enrique Tejedo went from 77 of 98 to one of 100. Alex Figueiras had a more similar presence: 54 of 98 and 62 of 100.

THE RESULT DESCRIBES THESE TWO CONFIGURATIONS

The browser model is unknown and the search and presentation conditions differ. We cannot attribute the differences solely to the interface or assume that any other API configuration would produce this result.

04 / WHERE EACH CANDIDATE APPEARS

Appearing often and appearing near the top

Being in a list and occupying its first positions describe different aspects of visibility. Position refers to the order of ChatGPT's text, from first to fifth.

Candidate	First Browser	First API	Top 3 Browser	Top 3 API
Enrique Tejedo	26	1	66	1
David Cerdá (Seomar)	24	6	45	29
Alex Figueiras	25	2	39	27
Jesús Porres	10	50	31	74
Queseo	0	28	10	77
Rocío Santamaría	0	0	17	27

These are counts of answers, based on 98 for the browser and 100 for the API. «Top 3» includes first, second and third position.

Two examples that change the interpretation

Queseo was the most repeated brand in the API and was in the top three in 77 answers. Jesús Porres topped more lists: 50 versus 28. In the browser, Queseo appeared 37 times, ten in the top three and none in first place.

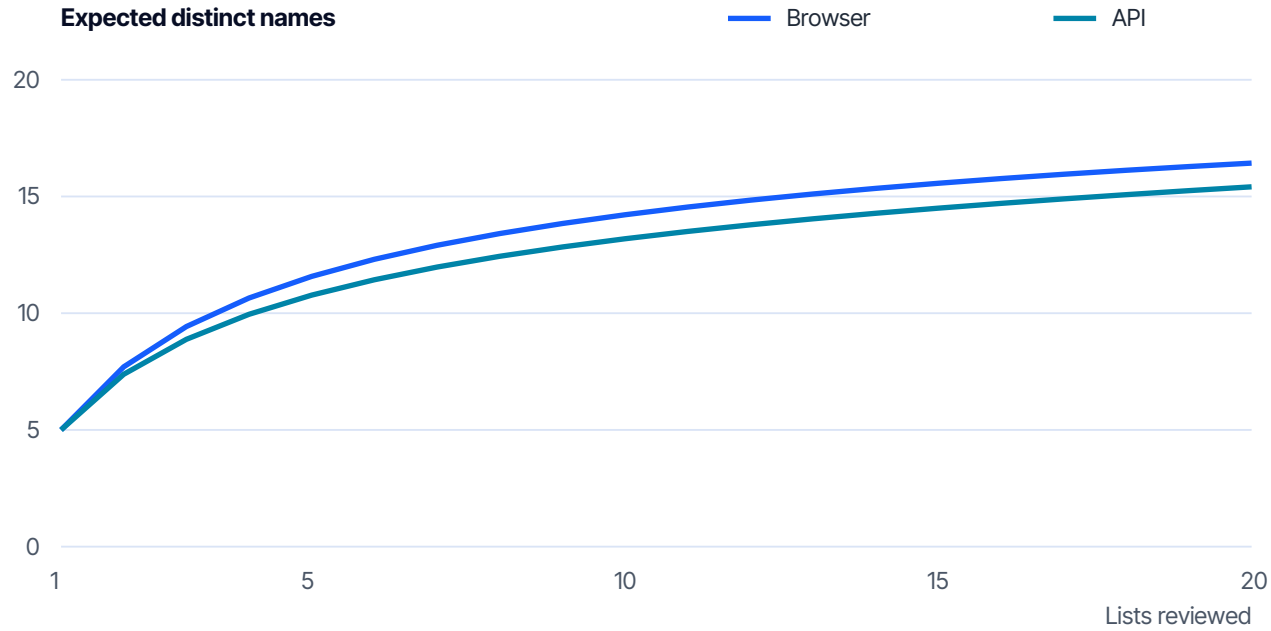
Alex Figueiras appeared in 54 browser lists and 62 API lists, but topped 25 of the former and only two of the latter. A similar frequency can accompany very different positions.

First place was shared among seven candidates in the browser and nine in the API. None of these measures alone assesses the professional quality of those recommended.

05 / WHAT ONE LIST LEAVES OUT

Five options do not exhaust the pool

The 198 lists contain **26 different candidates**: 17 appear in both environments, three only in the browser and six only in the API. Reviewing more answers reveals options missing from the first.



Lists reviewed	1	3	10	20
Expected candidates · browser	5.0	9.4	14.2	16.4
Expected candidates · API	5.0	8.9	13.2	15.4

This is a calculation on the saved lists: answers are randomly selected without replacement and the average number of distinct candidates is calculated. It does not predict how many new names ChatGPT will offer in the future.

A broad pool with concentrated presence

The nine most frequent candidates account for 83.5% of the 490 browser positions. In the API, the eight most frequent account for 82.6% of the 500 positions. Even so, the combination of each environment's five most frequent candidates appears only three times in the browser and five in the API.

06 / WHICH NAMES APPEAR TOGETHER

There are patterns within the variation

Lists change, but some pairs of candidates tend to appear together more often than others. These examples come from the 98 browser lists.

Two candidates	Lists in which they appear together
Queseo and Rocío Santamaría	24
Alex Figueiras and Álvaro Pichó	29
Paco Rubén and Alfredo Río	10
Paco Rubén and Alex Figueiras	0

For example, Queseo appears in 37 lists and Rocío Santamaría in 35. They coincide in 24. The API has a different distribution: Queseo and Jesús Porres coincide in 72 of its 100 lists.

The logged sources also vary

The 97 browser answers retaining a site log contain 80 different sets of domains. Ranking answer pairs by the similarity of those sets, the third with the most similar sources shares about 2.8 candidates; the least similar third shares about 1.8.

A RELATIONSHIP THAT HELPS INVESTIGATION

Candidate selection and logged sources are related. A search may supply names, and a name already selected may trigger another search to complete its profile. The available log does not reconstruct the full sequence or establish a cause in each case.

07 / WHAT THE ANSWERS LOOK LIKE

A stable format with changing names

BROWSER

582 words

Median: the middle answer when the 98 are ranked by length.

API

405 words

Median of the 100 API answers, calculated using the same criterion.

Observed element	Browser	API
Comparison table	98 of 98	0 of 100
Image elements in raw text	72 of 97	No equivalent visual record
Visible map	32 of 98	No equivalent visual record
Literal phrase «no es un ranking»	23 of 98	86 of 100

The browser usually presents criteria, descriptions of the five candidates, a comparison and suggestions based on the project. The API offers a shorter list with links and guidance on choosing.

The count of «no es un ranking» (not a ranking) is literal, after removing formatting and normalising spaces. Other answers say «no strict order of preference» or that there is no best professional for everyone. Those variants are not included in the row above.

Length is approximate: text is counted after removing links and internal markers. The mean, which very long or short answers can shift, is about 575 words in the browser and 400 in the API. Mean and median are different measures.

08 / WHY SOMEONE IS RECOMMENDED

Specialties are part of the message

Each list includes five explanations. Their body text was analysed without counting the heading name or recap tables again. These are mentions of topics, not checks that a service or credential is genuine.

Topic in the description	Browser 490 profiles	API 500 profiles
Local SEO: searches in an area	259 (53 %)	234 (47 %)
Technical SEO: how the website works	220 (45 %)	225 (45 %)
Online shops and ecommerce	179 (37 %)	227 (45 %)
AI, GEO and generative search	204 (42 %)	255 (51 %)
Years of activity or active «since» a year	224 (46 %)	216 (43 %)
Specific number of reviews or opinions	128 (26 %)	36 (7 %)
Success stories, testimonials or portfolio	16 (3 %)	42 (8 %)

Categories can overlap within a description. The same search vocabulary was applied to profile body text in both environments. Percentages are based on the 490 browser descriptions or 500 API descriptions. Counts are based on textual mentions.

What this analysis adds to GEO

Visibility goes beyond a name appearing. The need ChatGPT associates it with also matters: a migration, a local business, an online shop or a multilingual project. For a brand, reviewing the full description helps explain the proposition the reader receives.

09 / THE RECURRING PROFILE

Each brand retains its own associations

Although a brand enters and leaves the lists, certain topics recur in its descriptions. These examples help interpret the content of recommendations as well as count appearances.

Candidate	Association observed in the texts
Queseo	Migrations in all 37 browser descriptions and 87 API descriptions. Accompanied by technical SEO, online shops or international projects.
Rocío Santamaría	AI, GEO or generative search in 35 of 35 browser descriptions and 40 of 46 API descriptions. The year 2017 appears in 33 of the 35 browser descriptions.
Jesús Porres	Local businesses and professional services recur, such as clinics, dentists, lawyers and advisory firms.
Jamil Mansilla	Comprehensive work is described: analysis, content, optimisation and monitoring. The first API answer mentions 14 reviews.

How claims are presented

Sometimes the assistant states a specialty directly. At other times it attributes the information: «states», «according to their website» or «publishes». For example, API answer 48 attributes more than 16 years of experience to Queseo. Attribution tells us who makes a claim; it does not verify it.

A FAVOURABLE TONE IS NOT AN INDEPENDENT ASSESSMENT

The question asks for recommendations, so it encourages favourable explanations. The texts include cautions and advice before hiring, but this sample does not compare professional results or audit candidates' success stories.

10 / WHERE THE INFORMATION COMES FROM

Four things to distinguish

Measure	What it means here
Logged site or page	The domain or address appears in the saved log for that answer.
Citation or reference	The assistant identifies a source alongside a claim. Some internal labels do not by themselves resolve the exact URL.
Direct link	An address the reader can open from the answer.
Recommendation	The candidate occupies one of the five positions in the main list.

IP10 shows why these are different measures

The domain **ip10.es** appears in 96 of the 97 browser site logs. Its Valencia comparison appears in all 97 page logs. The source label «IP10» is in 68 of the 98 visible texts. A direct link to the domain appears in three of the 97 raw texts.

However, **IP10 holds no position in any of the 198 candidate lists**. Publishing a page that appears in the logs and being recommended as a provider are different outcomes.

WHAT CAN BE COMPARED WITH THE API

The API retained links, but not the complete sources for each search. IP10 receiving links in four API answers does not mean it was consulted only four times. OpenAI's documentation distinguishes citations from the full list of sources. [5]

The browser's page column also includes destinations linked in the text. It is used as an inventory of addresses, not as evidence that all were read before selecting the candidates.

11 / BROWSER SITE LOGS

A website may appear without being recommended

The column «Sitios consultados (1ª)» (sites consulted, first question) retains 97 logs. Each contains seven to thirteen domains, averaging 10.06. The five most frequent domains are ip10.es (96), valencianamente.es (90), albertofuster.com (79), sergidoseo.com (72) and jesusporres.com (68).

Candidate	Appears when their site is logged	Appears when their site is not logged
Álvaro Pichó	33 of 35 (94.3%)	3 of 62 (4.8%)
Alex Figueiras	44 of 56 (78.6%)	10 of 41 (24.4%)
Jesús Porres	46 of 68 (67.6%)	1 of 29 (3.4%)
Paul Cris	21 of 46 (45.7%)	1 of 51 (2.0%)
Andrés Castillo	16 of 51 (31.4%)	0 of 46 (0.0%)

For each candidate, the table divides the 97 answers into two groups. For example, Álvaro Pichó's site is logged in 35: he is recommended in 33 of them. Of the other 62, he appears in three.

The relationship does not apply equally to all

Queseo and Rocío Santamaría appear frequently although their domains hardly feature in this column. They may be described in comparisons or other sources, but each appearance cannot be attributed to a specific page. Alberto Fuster is the reverse case: his domain is logged 79 times and he appears as a candidate in only two of those answers.

For each answer, we do not know what information was processed before or after selecting the names. These relationships are descriptive and do not show that including a page among the sources causes a recommendation.

SEVERAL WEBSITES, ONE CANDIDATE

Enrique Tejedo appears in **77 of 98 browser lists (78.6%)** and one of 100 API lists. His description links to **kenjimarketing.es in 44 answers** and **seotejedo.es in two**; both sites are linked to his activity. Measuring only seotejedo.es would hide part of his link presence. Several brands and websites may offer more ways to identify the same professional: a plausible hypothesis, but this test does not show that adding domains causes his 77 recommendations. See the domain details.

12 / WHERE THE READER CAN GO

A link does not always identify a claim's source

LINKS IN API DESCRIPTIONS

500

One per candidate: five in each of the 100 answers.

DESTINATION: CANDIDATE'S DOMAIN

478 of 500

95.6% point to a domain identified with the recommended candidate.

Linked domain	Browser 97 raw texts	API 100 answers
queseo.es	32	82
jesusporres.com	42	82
alexfigueirasseo.com	39	62
rafasospedra.com	11	40
seomar.es	70	40
rociosantamaria.com	35	39
jamilmansilla.com	7	22

Each cell counts answers with at least one direct link to the domain, not visits received or searches performed. The other 22 API links point to third parties: ten go to IP10, spread across four answers.

The API includes the tag **utm_source=openai** in 408 of its 500 links. The tag may help identify clicks if analytics capture them, but these answers include no traffic or conversion data.

13 / LISTINGS AND MAPS

Two different signals in the browser

AT LEAST ONE LISTING LOGGED

40 of 98

The Excel file saves a business listing with details such as rating, reviews or website.

VISIBLE MAP

32 of 98

The answer includes a visible map: 31 in raw text and one in visible text.

Both signals coincide in 31 answers. Nine have logged listings without a visible map, and one has a visible map without listings in the log. It is therefore incorrect to say that 40 maps were shown.

Candidate	With listing 40 answers	Without listing 58 answers
Enrique Tejedo	37 (92.5 %)	40 (69.0 %)
Alfredo Río	23 (57.5 %)	6 (10.3 %)
Álvaro Pichó	3 (7.5 %)	33 (56.9 %)
Jesús Porres	11 (27.5 %)	37 (63.8 %)

The distribution changes between these groups. This suggests listings should also be recorded when measuring visibility. The association does not prove that displaying a listing adds or removes a candidate: both may be part of the same selection.

The saved API data contain text and links. They provide no equivalent visual observation for directly comparing the presence of maps.

14 / CHECKING DATA AND IDENTITY

A name, a listing and a link may disagree

The supplied resources reveal associations that should be reviewed. The original ChatGPT text, saved listings and Excel extraction columns must be distinguished.

Example	What the file contains
Jamil Mansilla	The saved listing associates his name with cristaleriasmansilla.com in five answers. In answer 2, the link in the main text leads to jamilmansilla.com.
Rocío Santamaría	The listings in answers 75 and 80 associate her name with dramelguizo.com and 229 reviews. The text of answer 80 links to rociosantamaria.com and mentions 22 reviews.
Jesús Porres	The listing log associates his name with jesteve.com in three answers. This logged relationship requires checks of identity and provenance.

These discrepancies indicate a mismatch between fields. The available files do not establish in every case whether it arose in the interface, listing selection or extraction. Nor do they prove a commercial relationship between the businesses.

Reviews: the same number, an unconfirmed source

The first API answer attributes **22 reviews to Rocío Santamaría** and **14 to Jamil Mansilla**, linking to their respective websites. Both numbers match the IP10 comparison consulted for this review. The match is consistent with that provenance, but does not prove it.

The external page reviewed now cannot replace a copy from the test date. Differences in reviews or years of activity may reflect different dates, sources or entities. No overall error rate has been calculated.

15 / HOW MANY ANSWERS TO COLLECT · 1 OF 3

Counting appearances and estimating a frequency

To measure a brand's presence, we count how many lists include it and divide by the valid lists in that environment. Each list counts once, even if the text mentions the brand several times.

QUESEO IN THE BROWSER

37.8 %

Appears in 37 of the 98 valid lists from free ChatGPT in the browser, without signing in.

QUESEO IN THE API

87.0 %

Appears in 87 of the 100 valid API lists with web search.

What we know from these data

These percentages describe the saved answers exactly. Estimating how often the brand would appear if the test were repeated under the same conditions involves uncertainty: another batch might give a different percentage.

Why a small batch is more sensitive

This is a fictional example. In both cases the brand appears in 40% of lists. Imagine one answer changes to include it: the percentage moves much more when there are few lists.

Hypothetical example	Batch of 10 lists	Batch of 100 lists
Before the change	4 of 10 = 40%	40 of 100 = 40%
One answer starts including the brand	5 of 10 = 50%	41 of 100 = 41%
Change in the percentage	+10 points	+1 point

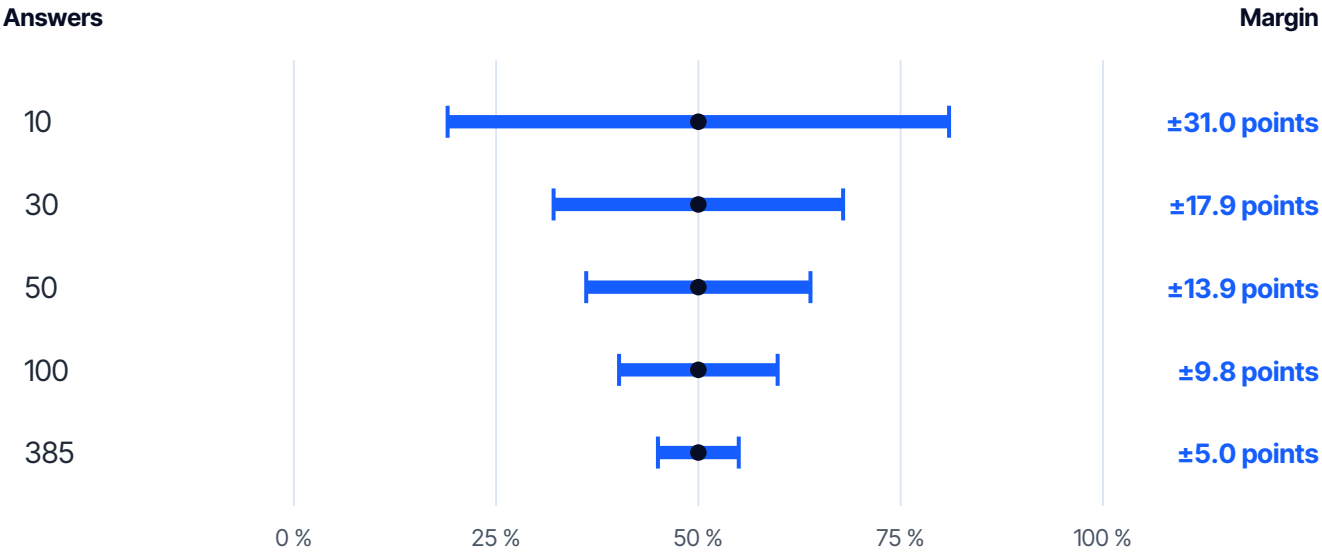
WHAT «PERCENTAGE POINTS» MEANS

It is the difference between two percentages. Going from 40% to 50% is an increase of 10 percentage points. On the following pages, «±9.8 points» means adding and subtracting 9.8 from the reference percentage.

15 / HOW MANY ANSWERS TO COLLECT · 2 OF 3

How to read a margin of ±9.8 points

The margin helps plan how many answers to collect. This hypothetical example starts at 50% presence. Each line shows 50% plus and minus the calculated margin. As more answers are collected, the range narrows.



An example with 100 answers

With 100 answers, the reference margin is **±9.8 points**. Applied to the example's 50%, the range is approximately **40.2% to 59.8%**: 50 minus 9.8 and 50 plus 9.8. With 385 answers it falls to about ±5 points: 45% to 55%.

Why 50% is used

50% is the case with the greatest variation between appearing and not appearing. It is used as a cautious planning reference. For a specific brand, the margin also depends on its number of appearances, as the next page shows.

HOW TO USE THESE FIGURES

Sample sizes apply to each environment separately. Greater precision requires more answers collected under comparable conditions. The chart is a 95% statistical approximation, especially tentative for small batches; ChatGPT does not guarantee these margins.

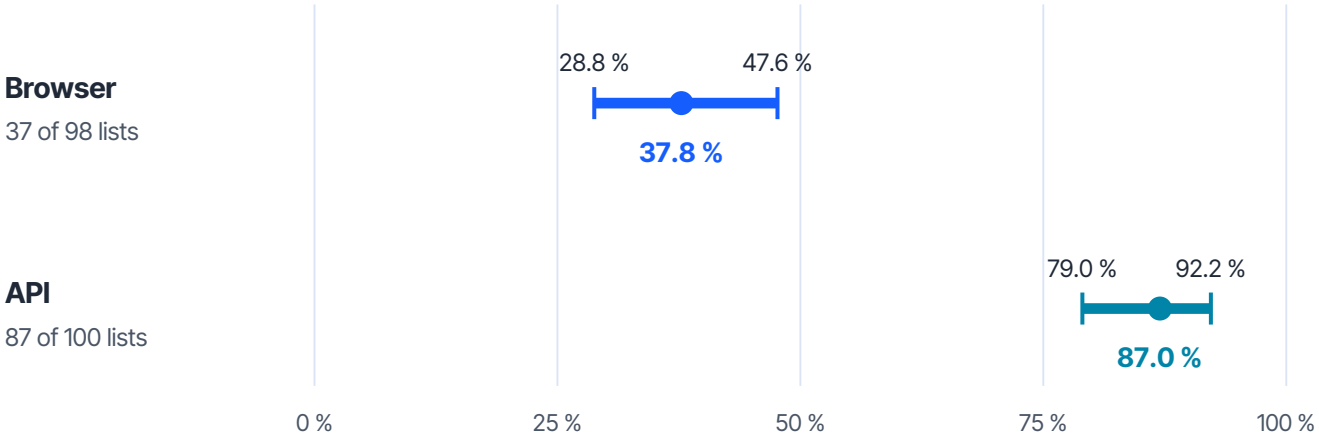
Calculation: 1.96 × square root of (0.25 / n), with n answers; multiply by 100 to express it in points. Assumes independence and a stable frequency, as explained next. [6]

15 / HOW MANY ANSWERS TO COLLECT · 3 OF 3

The interval applied to a brand in the sample

The chart uses Queseo's actual data. The point is the observed percentage; the line is the 95% confidence interval, calculated using the Wilson method. This interval summarises uncertainty when estimating a stable frequency from a limited batch. [6]

Queseo's presence: observed point and indicative 95% interval



How to read the browser result

37.8% is the result in the 98 saved lists. The range from **28.8% to 47.6%** expresses the estimate's imprecision under the assumptions below. In the API, the observed figure is **87.0%** and the range is **79.0% to 92.2%**.

What «95%» means

Imagine repeating the test many times with an appearance frequency that stays stable. A different interval would be calculated for each batch. The method is designed so that, under its assumptions, about 95 of every 100 intervals contain that frequency. The 95% refers to the calculation method.

THE TWO ASSUMPTIONS OF THE CALCULATION

Independence: one answer does not influence the next.
Stability: the appearance frequency stays stable during measurement, with the same question and comparable conditions.
This test does not establish that both assumptions hold.

These ranges do not predict the next answer or next week. Nor do a hundred queries from one night by themselves represent all users, dates, models or locations. See the intervals for all 26 candidates.

16 / HOW TO USE THE RESULTS

GEO monitoring that can be interpreted

1. Define the use you want to observe

Keep the exact question and specify the environment: in this study, free ChatGPT in the browser without signing in, or a specific API configuration.

2. Repeat and retain failures too

Count attempts and valid lists separately. Use an initial batch to understand variation, then increase it according to the precision needed. This test establishes no universal minimum.

3. Measure presence, position and links

Report how many lists include the brand, how many place it in the top three and how many let readers reach its website. Always explain the denominator.

4. Review the recommendation’s content

Record the service attributed, the client type it is suggested for and the specific details given. Check important claims against their source.

5. Log sources, listings and maps separately

Retain the available data. A link to a candidate’s website does not by itself reveal where its review count came from.

6. Compare dates and equivalent conditions

Alternate batches between environments and repeat them on different days. A change in model or configuration must be documented.

APPENDIX A / THE 26 CANDIDATES AND THEIR POSITIONS

Browser results

Free ChatGPT in the browser, without signing in. Base: 98 valid lists.

Candidate	Lists	Presence	1st	2nd	3rd	4th	5th
Enrique Tejedo	77	78.6 %	26	26	14	7	4
David Cerdá (Seomar)	71	72.4 %	24	12	9	16	10
Alex Figueiras	54	55.1 %	25	9	5	9	6
Jesús Porres	48	49.0 %	10	6	15	6	11
Queseo	37	37.8 %	0	5	5	13	14
Álvaro Pichó (SEO Levante)	36	36.7 %	9	17	0	7	3
Rocío Santamaría	35	35.7 %	0	3	14	7	11
Alfredo Río (SEO Local Valencia)	29	29.6 %	0	9	6	5	9
Paul Cris (Paul Marginean)	22	22.4 %	0	2	10	2	8
Andrés Castillo (GEO Valencia)	16	16.3 %	3	1	5	4	3
Paco Rubén	16	16.3 %	0	3	7	3	3
Rafa Sospedra	14	14.3 %	0	2	5	5	2
Jamil Mansilla	8	8.2 %	0	0	0	4	4
Daimako	7	7.1 %	0	0	0	5	2
Ruk Lab	7	7.1 %	0	0	1	4	2
Softalian (Nacho Monterde)	6	6.1 %	0	2	2	0	2
agenciaSEO.eu	3	3.1 %	0	0	0	1	2
Alberto Fuster	2	2.0 %	1	1	0	0	0
Alberto Illobre	1	1.0 %	0	0	0	0	1
Miguel Valero	1	1.0 %	0	0	0	0	1
Alban Febrer	0	0.0 %	0	0	0	0	0
Arturo Argilés (SEO&SEM Valencia)	0	0.0 %	0	0	0	0	0
Idital	0	0.0 %	0	0	0	0	0
Jorge Segovia	0	0.0 %	0	0	0	0	0
SEO Solutions	0	0.0 %	0	0	0	0	0
Vicente Climent (Santi Pc)	0	0.0 %	0	0	0	0	0
Total positions	490		98	98	98	98	98

Descending order by number of lists; ties by name. «Presence» is the percentage of valid lists including the candidate. Positions sum their appearances. Zero means they did not appear in this environment's sample.

See the API table

APPENDIX A / THE 26 CANDIDATES AND THEIR POSITIONS

API results

OpenAI API with web search. Base: 100 valid lists.

Candidate	Lists	Presence	1st	2nd	3rd	4th	5th
Queseo	87	87.0 %	28	37	12	5	5
Jesús Porres	82	82.0 %	50	15	9	7	1
Alex Figueiras	62	62.0 %	2	14	11	19	16
Rocío Santamaría	46	46.0 %	0	10	17	16	3
David Cerdá (Seomar)	44	44.0 %	6	7	16	8	7
Rafa Sospedra	42	42.0 %	10	7	8	10	7
Miguel Valero	26	26.0 %	1	6	7	6	6
Andrés Castillo (GEO Valencia)	24	24.0 %	0	1	7	3	13
Jamil Mansilla	23	23.0 %	1	0	2	7	13
Paco Rubén	17	17.0 %	0	0	0	6	11
Álvaro Pichó (SEO Levante)	12	12.0 %	1	3	5	2	1
Jorge Segovia	11	11.0 %	0	0	0	2	9
Softalian (Nacho Monterde)	9	9.0 %	0	0	4	4	1
Alberto Fuster	4	4.0 %	0	0	1	1	2
Alban Febrer	3	3.0 %	0	0	0	1	2
Alfredo Río (SEO Local Valencia)	1	1.0 %	0	0	1	0	0
Arturo Argilés (SEO&SEM Valencia)	1	1.0 %	0	0	0	0	1
Enrique Tejedo	1	1.0 %	1	0	0	0	0
Idital	1	1.0 %	0	0	0	1	0
Paul Cris (Paul Marginean)	1	1.0 %	0	0	0	1	0
Ruk Lab	1	1.0 %	0	0	0	0	1
SEO Solutions	1	1.0 %	0	0	0	0	1
Vicente Climent (Santi Pc)	1	1.0 %	0	0	0	1	0
agenciaSEO.eu	0	0.0 %	0	0	0	0	0
Alberto Illobre	0	0.0 %	0	0	0	0	0
Daimako	0	0.0 %	0	0	0	0	0
Total positions	500		100	100	100	100	100

Descending order by number of lists; ties by name. «Presence» is the percentage of valid lists including the candidate. Positions sum their appearances. Zero means they did not appear in this environment's sample.

See the browser table

APPENDIX B / HOW SIMILAR ANSWERS ARE · 1 OF 4

How two lists are compared

A **pair** consists of two lists. Each list is compared with every other list in its environment. In the **cross comparison**, each browser list is compared with each API list. This summarises how similar their recommendations are.

Measure	Browser	API	Cross comparison
Answer pairs compared	4,753	4,950	9,800
Mean shared names	2.30	2.62	1.77
Mean matches in name and position	0.58	0.81	0.32
Pairs with the same five names	47	59	10
Pairs with identical ordered lists	5	5	0
Pairs with the same first name	985	1,664	729

Reading the measures. Two lists share a name if they include the same candidate, wherever placed. Matching in **name and position** requires the same position in both. An identical list retains all five names and their order.

Browser example. There are 47 pairs with the same five names. Only five of those pairs also retain the order: 5 of 4,753, about 0.1% of all pairs. These are pairs of answers, not 47 or five new queries.

Distribution of shared names

Names in common	0	1	2	3	4	5
Browser	207	991	1,479	1,382	647	47
API	82	516	1,530	1,972	791	59
Cross comparison	1,187	2,546	3,783	1,938	336	10

For example, 1,479 in the «Browser» row, column «2», means that 1,479 pairs share exactly two names. Each row's columns sum to all pairs compared in that case.

WHY THOUSANDS OF PAIRS COME FROM JUST 198 LISTS

The same list is compared with many others. The 98 browser lists form 4,753 distinct pairs. These pairs reuse the same answers and do not amount to thousands of independent queries.

APPENDIX B / HOW SIMILAR ANSWERS ARE · 2 OF 4

Repeating names and retaining their order

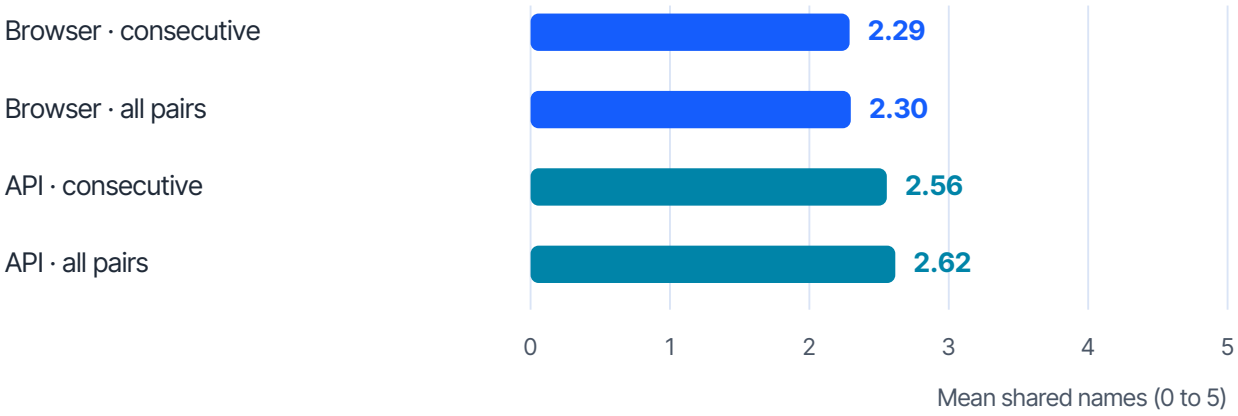
A recommendation can repeat in two ways. A **combination** repeats the same five candidates, even if positions change. An **ordered list** repeats all five candidates in exactly the same positions.

The most frequent repetition in each environment	Browser	API
Same five candidates, in any order	5 answers	6 answers
Same five candidates in the same order	3 answers	2 answers

Among the 98 browser lists, the most repeated combination contains the same five candidates in five answers. Five is the maximum for one specific combination; other combinations may also repeat.

Are consecutive answers more similar?

Each valid list is compared with the next: first with second, second with third, and so on. The average is compared with that of all pairs in the same environment, including those farther apart within the batch.



HOW TO READ THE CHART

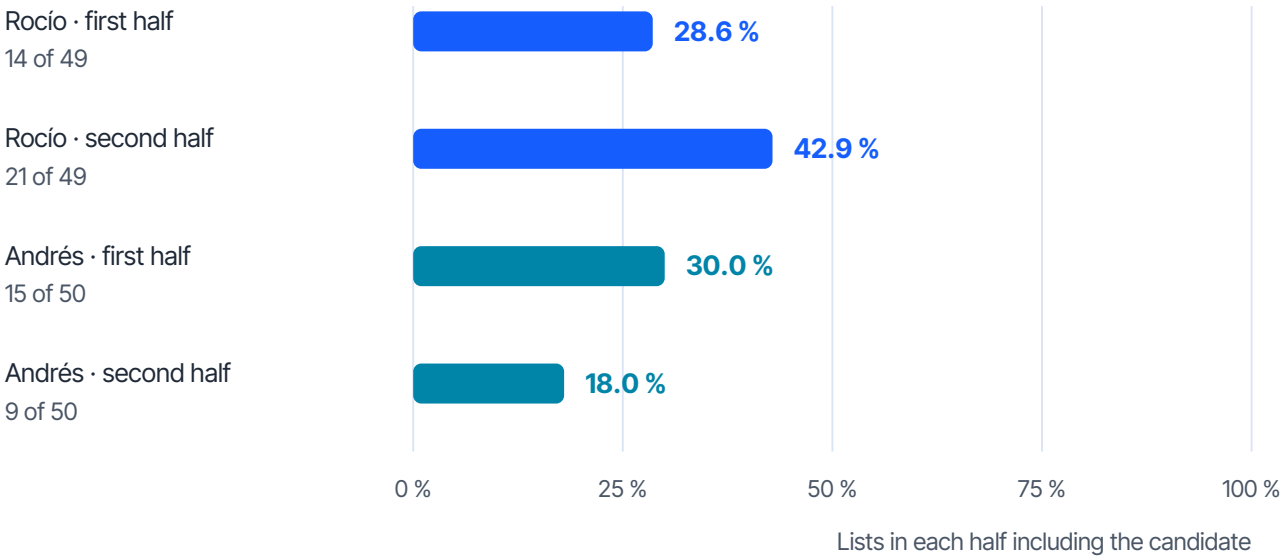
Out of five names, consecutive browser lists share an average of 2.29; across all pairs, 2.30. The API figures are 2.56 and 2.62. These averages are close: proximity within this batch barely changes this measure.

The comparison follows the order of valid lists; it does not set an identical time gap between queries. Similar averages neither establish independence of answers nor rule out other changes during measurement.

APPENDIX B / HOW SIMILAR ANSWERS ARE · 3 OF 4

What changes from start to finish

To see whether candidate presence changes during the batch, the valid lists in each environment are split into two halves, preserving order. The browser has **49 lists and 49 lists**; the API, **50 and 50**. Percentages are calculated within each half.



Browser: Rocío Santamaría's example

Rocío appears in 14 of the first 49 lists (**28.6%**) and 21 of the last 49 (**42.9%**). The difference is **+14.3 percentage points**. This is the largest absolute change between halves observed among browser candidates.

API: the largest change is 12 points

Andrés Castillo (GEO Valencia) goes from 15 of 50 lists (**30.0%**) to 9 of 50 (**18.0%**): **-12.0 points**. This shows the largest change in size, upward or downward; it is not the average change across all candidates.

WHAT THIS COMPARISON ESTABLISHES

Observed presence varies between the two parts of the batch. This split cannot separate variation between answers from a possible system change. Nor does it establish a trend that will continue at other times or dates.

APPENDIX B / HOW SIMILAR ANSWERS ARE · 4 OF 4

Similarity and discovery of new names

Jaccard: how many names two lists share

This index divides shared names by all distinct names in the combined lists. It ranges from **0** (no shared names) to **1** (the same five names). Candidate order is not part of the calculation.

Fictional example	Five candidates in each list
List A	Queseo · Rocío Santamaría · Jamil Mansilla · Alex Figueiras · Jesús Porres
List B	Rocío Santamaría · Queseo · Jamil Mansilla · Enrique Tejedo · Rafa Sospedra

The lists share **three names**: Queseo, Rocío and Jamil. Combined, they contain **seven different candidates**. Therefore, Jaccard = 3 / 7 = **0.429**. Each person or brand counts once in the set of distinct names.

Average in the actual data	Browser	API	Cross comparison
Jaccard index	0.326	0.378	0.234

The higher the index, the more similar the sets of names. The browser's 0.326 is an average pairwise similarity; it does not mean that 32.6% of its answers are identical.

Accumulation: how many distinct candidates are found

The accumulation curve answers another question: as more lists are reviewed, how many different names are collected? The first list gives five. Adding another does not count repeated names again. For example, selecting three browser lists gives an expected 9.4 distinct candidates, averaged over possible selections from the saved lists.

The calculation selects lists from this sample without selecting the same answer again within one selection. It describes diversity in the file, not how many names future queries will bring.

Details of the curve calculation. For each candidate, calculate 1 - C(N - f, k) / C(N, k): the probability of including them when choosing k lists. N is the total saved lists, f the lists including that candidate, and C counts ways of selecting lists without replacement. The curve sums these values over all candidates.

APPENDIX C / UNCERTAINTY UNDER STATISTICAL ASSUMPTIONS

Indicative intervals by candidate

Each range expresses presence as a percentage. «95% CI» means 95% confidence interval. Wilson is used assuming independence and stable conditions, neither demonstrated here. «Diff.» is API minus browser, in percentage points. [6]

Candidate	Browser · 95% CI	API · 95% CI	Diff. points
Jesús Porres	39.3 to 58.7	73.3 to 88.3	+33.0
Queseo	28.8 to 47.6	79.0 to 92.2	+49.2
Alex Figueiras	45.2 to 64.6	52.2 to 70.9	+6.9
David Cerdá (Seomar)	62.9 to 80.3	34.7 to 53.8	-28.4
Rocío Santamaría	26.9 to 45.6	36.6 to 55.7	+10.3
Enrique Tejedo	69.5 to 85.5	0.2 to 5.4	-77.6
Rafa Sospedra	8.7 to 22.6	32.8 to 51.8	+27.7
Álvaro Pichó (SEO Levante)	27.9 to 46.6	7.0 to 19.8	-24.7
Andrés Castillo (GEO Valencia)	10.3 to 24.9	16.7 to 33.2	+7.7
Paco Rubén	10.3 to 24.9	10.9 to 25.5	+0.7
Jamil Mansilla	4.2 to 15.3	15.8 to 32.2	+14.8
Alfredo Río (SEO Local Valencia)	21.5 to 39.3	0.2 to 5.4	-28.6
Miguel Valero	0.2 to 5.6	18.4 to 35.4	+25.0
Paul Cris (Paul Marginean)	15.3 to 31.7	0.2 to 5.4	-21.4
Softalian (Nacho Monterde)	2.8 to 12.7	4.8 to 16.2	+2.9
Jorge Segovia	0.0 to 3.8	6.3 to 18.6	+11.0
Ruk Lab	3.5 to 14.0	0.2 to 5.4	-6.1
Daimako	3.5 to 14.0	0.0 to 3.7	-7.1
Alberto Fuster	0.6 to 7.1	1.6 to 9.8	+2.0
agenciaSEO.eu	1.0 to 8.6	0.0 to 3.7	-3.1
Alban Febrer	0.0 to 3.8	1.0 to 8.5	+3.0
Alberto Illobre	0.2 to 5.6	0.0 to 3.7	-1.0
Arturo Argilés (SEO&SEM Valencia)	0.0 to 3.8	0.2 to 5.4	+1.0
Idital	0.0 to 3.8	0.2 to 5.4	+1.0
SEO Solutions	0.0 to 3.8	0.2 to 5.4	+1.0
Vicente Climent (Santi Pc)	0.0 to 3.8	0.2 to 5.4	+1.0

An observed zero does not mean zero probability. In an exploratory analysis, Fisher compares frequencies across environments: 13 candidates yield $p < 0.05$; adjusting for 26 comparisons (Bonferroni) leaves ten. These tests require the assumptions above and neither prove causality nor support generalising to other conditions.

APPENDIX D / DOMAINS ASSOCIATED WITH NAMES · 1 OF 3

Which domains appear alongside each name

51 associations between the 26 candidates and domains linked in their descriptions or listings. They include sites identified with their activity, external sources and unconfirmed attributions.

Name or brand	Associated domain	Sites Browser	Link Browser	Link API	Listing Browser
Enrique Tejedo	kenjimarketing.es	0	44	0	0
Enrique Tejedo	seotejedo.es	2	2	1	36
Enrique Tejedo	sergidoseo.com *	72	5	0	0
Enrique Tejedo	seotejedo.com *	0	4	0	0
Enrique Tejedo	enriquetejedo.com *	0	3	0	0
Enrique Tejedo	google.com *	0	3	0	0
Enrique Tejedo	valenciatop.info *	27	1	0	0
David Cerdá (Seomar)	seomar.es	3	70	40	20
David Cerdá (Seomar)	ip10.es *	96	0	4	0
Alex Figueiras	alexfigueirasseo.com	56	39	62	0
Alex Figueiras	sergidoseo.com *	72	7	0	0
Jesús Porres	jesusporres.com	68	42	82	0
Jesús Porres	jesteve.com *	0	0	0	3
Queseo	queseo.es	1	32	82	6
Queseo	ip10.es *	96	1	3	0
Queseo	valencia.wordcamp.org *	0	0	2	0

Sites Browser: the domain is in the answer's site log, whether or not that candidate is recommended (97 logs). **Link Browser / API**: linked within that candidate's section (97 browser texts / 100 API texts). **Listing Browser**: appears as the website in a listing associated with that name (98 browser lists).

* External site, platform or address not treated as the candidate's confirmed own website. An association in an answer or listing establishes neither ownership nor that the link is correct. Columns may overlap and are not summed.

AN EXAMPLE FOR INTERPRETING THE COLUMNS

Kenji appears in 44 links in Enrique's section and in 0 site logs. Being linked in an answer does not show it was consulted before selecting the candidate.

APPENDIX D / DOMAINS ASSOCIATED WITH NAMES · 2 OF 3

Which domains appear alongside each name

Domain table continued. Each figure counts answers, and each domain counts at most once per column and answer.

Name or brand	Associated domain	Sites Browser	Link Browser	Link API	Listing Browser
Álvaro Pichó	seolevante.com	35	31	12	3
Álvaro Pichó	geovalencia.com *	51	1	0	0
Rocío Santamaría	rociosantamaria.com	0	35	39	6
Rocío Santamaría	malt.es *	28	0	4	0
Rocío Santamaría	ip10.es *	96	0	1	0
Rocío Santamaría	malt.com *	0	0	1	0
Rocío Santamaría	mujeresenseo.es *	0	0	1	0
Rocío Santamaría	dramelguizo.com *	0	0	0	2
Alfredo Río	seolocalvalencia.es	29	25	1	21
Paul Cris / Marginean	paulcris.com	46	18	1	4
Andrés Castillo	geovalencia.com	51	15	24	0
Paco Rubén	pacoruben.com	5	14	14	13
Paco Rubén	elpais.com *	44	0	1	0
Paco Rubén	ip10.es *	96	0	1	0
Paco Rubén	valencianamente.es *	90	0	1	0
Rafa Sospedra	rafasospedra.com	0	11	40	4
Rafa Sospedra	albertofuster.com *	79	0	1	0
Rafa Sospedra	malt.es *	28	0	1	0

Sites Browser: the domain is in the answer's site log, whether or not that candidate is recommended (97 logs). **Link Browser / API**: linked within that candidate's section (97 browser texts / 100 API texts). **Listing Browser**: appears as the website in a listing associated with that name (98 browser lists).

* External site, platform or address not treated as the candidate's confirmed own website. An association in an answer or listing establishes neither ownership nor that the link is correct. Columns may overlap and are not summed.

«Browser» means free ChatGPT in the browser, without signing in. The API uses web search. Links are counted within the main list, excluding those in the final summary.

APPENDIX D / DOMAINS ASSOCIATED WITH NAMES · 3 OF 3

Which domains appear alongside each name

Domain table continued. Each figure counts answers, and each domain counts at most once per column and answer.

Name or brand	Associated domain	Sites Browser	Link Browser	Link API	Listing Browser
Jamil Mansilla	jamilmansilla.com	0	7	22	0
Jamil Mansilla	ip10.es *	96	0	1	0
Jamil Mansilla	cristaleriasmansilla.com *	0	0	0	5
Daimako	daimako.com	2	7	0	6
Ruk Lab	ruklab.com	0	6	1	8
Softalian / N. Monterde	softalian.com	0	5	9	8
agenciaSEO.eu	agenciaseo.eu	0	3	0	0
Alberto Fuster	albertofuster.com	79	2	4	1
Alberto Illobre	valenciatop.info *	27	1	0	0
Alberto Illobre	sideor.com *	0	0	0	1
Miguel Valero	miguel-valero.es	1	1	26	0
Alban Febrer	albanfebrer.com	0	0	3	0
Arturo Argilés	seosemvalencia.es	1	0	1	0
Idital	idital.com	0	0	1	0
Jorge Segovia	jorgesegoviaciscar.com	0	0	11	0
SEO Solutions	seosolutions.es	0	0	1	0
Vicente Climent	santipc.com	0	0	1	0

Sites Browser: the domain is in the answer's site log, whether or not that candidate is recommended (97 logs). **Link Browser / API**: linked within that candidate's section (97 browser texts / 100 API texts). **Listing Browser**: appears as the website in a listing associated with that name (98 browser lists).

* External site, platform or address not treated as the candidate's confirmed own website. An association in an answer or listing establishes neither ownership nor that the link is correct. Columns may overlap and are not summed.

Counts come from the original Excel files. Domains are normalised by removing «www.», while «.es» and «.com» remain separate. Raw browser text is unavailable for one of the 98 lists, so links are calculated over 97.

SOURCES AND AUTHORSHIP

Documentation for interpreting the study

[1] resultados-consultas-navegador-chatgpt.xlsx.

[2] resultado-5.6-luna-100-format-consultas-api.xlsx.

[3] GEO framework. Aggarwal et al. (2024). GEO: Generative Engine Optimization.

[4] Verification of claims and citations. Liu, Zhang and Liang (2023). Evaluating Verifiability in Generative Search Engines.

[5] Search source logs. OpenAI. Web search, Sources section.

[6] Intervals for proportions. NIST/SEMATECH. Confidence intervals, section on the Wilson method.

Authorship and scope of the publication

Nina S. Dzhezher of ELNIQ SL (https://www.elniq.com/), and Eric Jorge Seguí Parejo of Queseo (https://queseo.es/). Queseo is also one of the brands recommended in the answers analysed. The same counting criterion is applied to all candidates.

Quantitative results come from the two original Excel files. External references provide methodological context; they do not expand the sample of answers.