

New IQRush Research as featured in SEJ article: Shows When AI Visibility Rankings Are Actually Ready to Trust

IQRush Research shows a simple test for telling a real change in AI search visibility from statistical noise & says AI visibility tools should Show their Math

SEATTLE, WA, UNITED STATES, July 15, 2026 /EINPresswire.com/ -- [IQRush](#), the verification layer for AI search visibility,

released new research today that answers a question most AI visibility vendors skip: how do you know when your numbers are ready to act on?



The paper, From Stochastic to Stable: Rank Stability and Structural Sufficiency in AI Visibility Measurement, is by IQRush co-founder and Chief AI and Data Officer Ron Sielinski. It is available now at [\[Link 2\]](#).

“

The industry moved fast to produce AI visibility scores and slower to ask whether those scores can support the decisions people make with them, our answer is simple to state: show your math.”

*Ron Sielinski, co-founder and
Chief AI and Data Officer at
IQRush*

The paper gives marketers a simple test for telling a real change in AI search visibility from statistical noise. Search Engine Journal covered the research, citing SparkToro founder Rand Fishkin's advice that before spending money on AI visibility tracking, buyers should make sure their provider “shows their math.” [\[Link 1\]](#)

AI answer engines like ChatGPT, Perplexity, and Gemini do not give the same answer twice. Ask the same question again and you can get a different set of cited sources. A citation score on your dashboard is one sample, not a fixed

fact. Two brands shown first and second on a leaderboard may really be tied.

Data quality has topped CMO priority lists for the past two years, and AI visibility is where that problem shows-up hardest. Most vendors point a tracker at an AI engine and report back the same way analytics always have run the query, count the citations, call it the score. That approach works for a system that gives the same answer every time. It does not work for one

that samples a different answer on every call. The dashboards are not unreliable because the tooling is bad. They are unreliable because a deterministic method is being used to measure a probabilistic system.

Sielinski tested 30 combinations of platforms and topics across Gemini, SearchGPT, and Perplexity. The number of citation-bearing answers needed before a ranking could be trusted ranged from 33 to 94. Three of the 30, all on SearchGPT, never reached that point even after 125 questions, because the top sources stayed too close together to tell apart.

The paper sets two tests a ranking must pass before anyone acts on it. First, has the order stopped changing as more data comes in. Second, are the gaps between brands bigger than the measurement's own margin of error. A ranking is only real when both are true.

"The industry moved fast to produce AI visibility scores and slower to ask whether those scores can support the decisions people make with them," said Ron Sielinski, co-founder and Chief AI and Data Officer at IQRush. "Our answer is simple to state: show your math. If a provider can't tell you whether a ranking has settled and whether the gaps between brands are real, the number isn't ready, no matter how clean the dashboard makes it look."

Sielinski presents related work this month at the IAB Measurement Leadership Summit in New York. The paper puts the underlying method on the record and makes the test repeatable for anyone measuring AI visibility, whether they run it in-house or buy it from a vendor.

Before trusting any AI visibility vendor or dashboard, ask:

- 1.How many times was each query run?
- 2.What is the margin of error on each score?
- 3.Would the ranking hold up if you ran it again next week?
- 4.Which positions are truly separate, and which are really tied?

IQRush built its platform to answer those four questions on every metric it reports. That is what decision-grade means.

Brands and agencies that want to see how their own numbers hold up under the same test can book a 20-minute walkthrough on the [IQRush.ai](https://iqrush.ai) site and you can find the paper Stochastic to Stable at the link below. .

About IQRush

IQRush is a precision measurement infrastructure for probabilistic AI models. The company

measures and verifies AI output before humans and AI agents make decisions, so they act on decision-grade data. IQRush reports every AI visibility metric, including citations and mentions across Google AI Overviews, Perplexity, ChatGPT, and Gemini, with the confidence intervals and readiness checks marketing and analytics teams need to make claims that hold up in front of a CMO or a board. IQRush is headquartered in Seattle, Washington. Learn more at iqrush.ai.

Resources :

From Stochastic to Stable (paper): <https://arxiv.org/html/2607.10341v1>

IQRush research and insights hub [Link 1] : <https://iqrush.ai/articles/how-to-tell-if-your-ai-search-visibility-number-is-real>

Search Engine Journal coverage [Link 2] : <https://www.searchenginejournal.com/ai-visibility-rankings-arent-stable-new-research-shows-its-mostly-statistical-noise/581905/>

Tracie Kambies

IQRush.ai

+1 678-427-4202

[email us here](#)

Visit us on social media:

[LinkedIn](#)

This press release can be viewed online at: <https://www.einpresswire.com/article/926904906>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.