

The Multilingual Advantage in AI Search

How multilingual press release distribution can increase AI citations



Foreword from Notified's CEO



“In an AI training set, silence is itself a signal. The brands that supply the content shape the answer.”

Erik Carlson

Chief Executive Officer,
Notified

Most of your audience today starts their search in AI, not by scrolling a page with ten blue links. AI inherently carries a bias most PR and communications professionals haven't even considered: it searches and answers in English, even when the person asking doesn't speak the language. This creates a real gap, and an opportunity.

If your brand only shows up in English, you can be invisible to the audiences you are working hardest to win in their own market. The model still answers, it just fills the response with whatever it can find, which is often a machine-translated guess about your headline, your brand, or your company.

We wanted to test whether distribution strategy could actually move the needle on this issue, so our team researched GlobeNewswire press release performance across 100 releases across 43 companies, building 50 matched pairs (English-only against translated releases), to measure real AI citations from first-party server data, not estimates.¹

The result surprised us.

The English version of a multilingual release earned an 82% median lift in AI citations and outperformed its English-only match in 92% of the pairs.¹ The translated versions generated much of their activity through indexing and training-related retrieval, while most citations went to the English version.

This is not an argument for translating every announcement. It is a case for translating important news when a real stakeholder audience exists in the target market and using citation data to measure the result.

The brands that win the next few years won't be the ones with the biggest communications budgets. They'll be the ones feeding the highest-quality, brand-identifiable, original content, in the right languages, into systems that increasingly translate your narrative on your behalf.

This is what this paper is all about. I think you'll find it useful. A special thanks to Jimmy Kropelin and Christian Eilert-Olsen, who led this research and authored the paper. This is exactly the kind of original, evidence-led thinking our industry needs right now, and it reflects the best of what the Notified team does.

Executive Summary

Buyers, journalists, and investors increasingly begin their research inside AI tools rather than a page of blue links. Those systems overwhelmingly source, weigh, and answer in English, even when the question is asked in another language. For global brands, that creates a visibility gap and a messaging risk at the same time.

To see if distribution strategy could close this gap, GlobeNewswire by Notified studied 100 press releases from 43 companies across 15 industries. We compared 50 pairs of similar releases from the same company and about the same topic. In each pair, one release was shared only in English, while the other was shared in several languages.


Using GlobeNewswire's own data, we tracked how often AI tools cited the English version of each release during the first 30 days after it was distributed.

The results were clear and statistically significant. When a press release was shared in several languages, its English version received a median increase of 82% in AI citations.

The multilingual release performed better in 46 of the 50 comparisons, with the largest increase reaching 1,064%. The results passed every test we used to check their strength. The chance that these results happened randomly was less than one in a million.¹

The data does not support translating every announcement into as many languages as possible. The lift concentrates where Multilingual distribution matches a real stakeholder audience. This paper lays out the finding, the mechanism behind it, and a practical framework for using citation data to inform global distribution.

The harder, more useful question is not whether more copies get more citations — it is whether professionally translated versions make the English release itself more citable. They do.



Discovery has a new gatekeeper

Buyers, journalists, analysts, and investors increasingly use AI tools to research companies and verify information. As these systems answer more questions directly, visibility depends partly on whether they can find, understand, and cite a company's content.

Press releases are well suited to this environment. They provide timely, structured, first-party information that AI systems can retrieve as source material. This means distribution decisions can influence more than immediate reach. They can also affect how often company information appears in AI-generated answers.

If AI systems are now a primary audience for corporate content, the strategic question is which distribution choices make content more likely to be retrieved and cited. This paper examines one of the most consequential and most overlooked of those choices: language.

Why language can affect AI visibility

Large language models have a language problem baked into their foundations. More than half of web content is in English, although only about 16% of the world speaks the language.² This imbalance can affect which sources AI systems retrieve and how they answer questions in different languages.³

- **AI tends to rely on English sources.** An analysis of more than 10 million ChatGPT searches found that a large share of the model's background research runs in English even when the user's query is not, and that a majority of non-English research sessions are supplemented with English-language sources.⁴ German users asking about German companies get shown international brands; local market leaders go missing. Brands without cross-language content are effectively invisible for those queries.
- **Multilingual systems need translation pairs.** Research on multilingual retrieval-augmented generation shows these systems perform best when multilingual sources are available, and identifies expert-verified translation pairs, an original document and a professionally translated counterpart, as critical for accuracy across languages.⁷
- **Machine translations can make contextual errors.** In adversarial tests, LLMs repeated or elaborated on false information in 50 to 82% of cases.⁶ When brands rely on automated translation, or leave translation to the model, inaccurate content

becomes a vector for misinformation about the brand. Stanford's Legal Design Lab documented real machine mistranslations where "charged with a battery case" became "loaded with a case of batteries" and "due date" became "date to give birth."⁷ Expert human translations provide the verified factual anchors that reduce this risk.

A study published in *Nature* found leading chatbots answer the same question differently depending on the language it is asked in, because each language corpus draws on different sources.⁸ The lesson is about supply. Whoever supplies the content in a language shapes the answer in that language. When your brand is absent, the model still answers, from whatever the corpus contains. In an AI training set, silence is itself a signal.

Together these findings pointed to a testable hypothesis: if AI systems treat cross-language corroboration as a signal of authority, then distributing a release in multiple languages should increase citations of the English version, not merely add foreign-language copies to the pile.

Measuring actual AI retrieval activity

Most of the AI visibility industry works in the theoretical. Prompt-tracking and AEO tools estimate visibility by running prompts and watching what comes back. They are useful for strategy, but they share the weakness of keyword tracking: if you miss a prompt, you miss the conversation.

GlobeNewswire by Notified took a different path. Through our integration with Profound, we use proprietary data to measure every time an AI crawler comes to a release to cite it, index it, or train on it.

This approach measures observed crawler activity on individual releases rather than estimating visibility from a predefined set of prompts. It allowed us to compare citation activity across the 50 matched pairs in this study.

Citation

The action

AI fetchers grab specific web pages in real time to fulfill specific user requests on AI assistants, such as answering a question or summarizing a webpage via an AI chatbot.

Indexing

The shortlist

AI search crawlers automatically scan and index websites for AI-powered search experiences.

Training

The muscle memory

When content is absorbed into the next model build.

Study design

We built the study to isolate the effect of multilingual distribution as cleanly as possible.



Two design choices carry the weight.

The matched-pair structure controls for company, brand strength, and content type. The comparison is always the same organization announcing the same kind of news.

We counted citations on the English version only. Counting foreign-language copies would guarantee a flattering result, because more copies of a release always produce more total citations. The harder and more useful question is whether the existence of professionally translated versions makes the English release itself more citable.


Parameter	Detail
Sample	100 press releases across 43 companies and 15 industries
Design	50 matched pairs: same company, comparable topic; one English-only and one Multilingual release per pair
Metric	AI citations on the English-language release only
Window	First 30 days post-distribution
Data source	First-party citation data from GlobeNewswire's Profound integration
Win rate	Instances when multi-language releases secured more citations than English-only releases

What we found


Across 50 matched pairs, the English version of a multilingual release received significantly more AI citations than its English-only match.



Composite visibility lift, English version (%)



Where each layer's activity comes from



The mechanism is consistent with the research in Section 2. Professionally translated versions may create stronger cross-language connections and additional signals that help AI systems cite the English release. The lift is not primarily about reaching foreign audiences, though that remains a real benefit. It is about how these systems weigh and source content.

The full footprint tells the deeper story. The English version alone shows a +82% citation lift, a +22% indexing lift, and a +96% training lift. But the translation layer's contribution is where the mechanism becomes visible: translations account for only 9% of the citations, yet 51% of the indexing and 54% of the training-data incorporation. When you look across all three layers, citation, indexing, and training, the composite visibility lift reaches +132%.

In other words, translated versions generated proportionally more indexing and training-related activity, while citations remained concentrated on the English version. This suggests that translated content may expand the release's cross-language presence even when the translated pages are not the sources ultimately cited.

The effect is broad, not a fluke of one segment. Thirteen of the fifteen industries in the sample showed 100% win rates (when multilingual distribution secured more citations than the English only release). B2B technology, the largest segment at 23 pairs, won 91% of the time with a median lift essentially identical to the rest of the sample, confirming the effect is not an artifact of any single industry. Public companies lifted slightly more than private ones (94% win rate and +109% median across 16 pairs). Given how many institutional investors now research through AI tools, publicly traded issuers arguably have the most direct financial stake in AI visibility of any audience segment.

Results also varied by announcement type.

Durable, audience-defining news lifts the most



Announcement Type	Win Rate	Median Lift
Regional expansion	+100%	+150%
Customer announcement	+100%	+123%
Report / study	+100%	+121%
Award	+80%	+106%
Product announcement	+92%	+93%
Partner announcement	+100%	+77%
Company announcement	+100%	+44%
Upcoming event	+67%	+19%



Regional expansion news, where a brand is by definition addressing a new-market audience, showed among the strongest lifts in the study. One possible explanation is that expansion news remains relevant longer, while interest in an event declines after it takes place.



What the data did not show

Rigorous research means being transparent about the limits. Two findings kept us honest, and they matter for how you spend a translation budget.

More languages did not mean more citations.

Across releases using one to nine additional languages, we found no correlation between language count and citation performance ($r \approx 0.05$). Releases with a single additional language sometimes outperformed releases with nine. The benefit activates with even one additional language. After that, chasing a citation multiplier does not pay.

More regions did not mean more citations.

Distribution geography showed essentially zero independent relationship with performance. It is the act of creating expert-translated content that drives the lift, not the breadth of the wire.

The takeaway translating just to translate does not work. We deliberately included releases from U.S.-only companies translating without a real audience in the target market, and their lift was minimal or slightly negative. The companies that won were the ones being researched by real stakeholders in the target language. The decision on how many languages to use should be driven by your stakeholder map.

Notably, 4 of the 50 pairs saw English-only out-cite Multilingual, and two of those four were announcement types that tend to generate fewer citations.

The 92% win rate is strong, but it does not definitively conclude that translation will improve results for every release.

Independent validation

Our finding does not stand alone. In March 2026, a separate team at Stacker and Scrunch examined a related question: whether distributing the same story across more publisher sites affects AI citations.

Their expanded study, ten times the sample of their December 2025 pilot, confirmed it does, with statistical significance: **a +239% median citation lift, a 97% citation rate, and roughly a threefold increase in cross-platform AI visibility.**⁹

The two studies measure different levers and point in the same direction. Stacker proved that broader earned distribution drives more citations. We proved that more expert-translated language versions drive more citations on the English version. A multilingual GlobeNewswire release creates both at once, new domain placements and new language entry points, and the two effects compound.

The investor and stakeholder imperative



For public companies, this has moved from a communications question to an investor relations one. Institutional investors are increasingly using AI tools as information-gathering instruments, and non-traditional media such as podcasts and industry newsletters now rival mainstream financial media in adoption among them.¹⁰ When half of institutional investors are querying AI tools, and those tools search in English first regardless of the investor's language, the brands feeding the highest-quality multilingual content into these systems will dominate the answers.

The risk runs the other way too. Many brands have non-English-speaking stakeholders, customers, investors, regulators, who can be handed an unintended narrative when the model fills a content gap with machine-translated guesswork. Three practical actions follow: translate your earnings transcripts, make sure your IR website clears the minimum bar for AI crawling, and use expert-translated press releases to keep your stakeholders accurately informed in their own language.

From data to strategy

The value of this study is not the instruction to go global. It shows that language and release type can be considered alongside audience needs when planning distribution. Four practical recommendations follow.



Prioritize the announcements that carry your narrative.

Partner, customer, company, and product news, and especially regional expansions, are where multilingual distribution reliably compounds. Durable news keeps earning citations long after publication.



Translate where your stakeholders are, not everywhere.

Start from your customer, partner, and investor footprint and choose languages to match it. One additional language, aimed at a real audience, is enough to activate the effect.



Use the wire as a feedback loop.

Because we pull actual citation counts on every release, you can see in real numbers whether a topic is resonating or sitting in an information desert, and validate or kill a theory before you invest more into it.



Extend your reporting window.

Most teams only look at results 24 hours after distribution. The real signal lives at 30 days and beyond. We are seeing views double over six months on some releases, and releases older than a year surge in citations. If you plan to expand into a market, getting accurate content into the training set months ahead of your go-to-market can pay dividends before you hire a single person there.

Every brand narrative now lives far beyond its publication date. AI systems retrieve, cite, and train on corporate content continuously, which makes the press release one of the most durable, structured, and machine-legible assets a brand controls. Distributing that asset in the languages your stakeholders speak is one of the few levers in communications with a measured, repeatable effect on AI visibility. Now it is also one you can verify.

LET'S FIND WHERE THIS PAYS OFF FOR YOU

See what's possible with multilingual releases

Bring us an upcoming announcement, and we'll show you how multilingual distribution can help expand reach, improve visibility, and connect with the audiences that matter most.

Get in touch



Methodology Notes

All 100 source releases were retained; no exclusions were made. Every matched pair contains exactly one English-only and one multilingual release from the same company on the same release topic. Citations were measured on the English-language version over a 30-day post-distribution window; citations from other languages were not counted. Time-order analysis confirms release sequencing does not explain the result. Company identities are anonymized. This is an observational matched-pair design, not a randomized experiment; companies choose which announcements to translate. The training-data metric is the most novel of the three, and the 30-day window captures only the near-term training signal; long-run behavior is observable across future model releases. The lift distribution is right-skewed, so medians lead every headline. Citation data is drawn from GlobeNewswire's first-party integration with Profound.

Sources

1. GlobeNewswire by Notified. (2026, May). *Global case study dataset: Multilingual press release performance and AI citations* [Unpublished internal dataset; first-party server-log measurement].
2. Center for Democracy & Technology. (2023). *Lost in translation: Large language models in non-English content analysis*.
3. ModernMT. (n.d.). *Making generative AI effectively multilingual at scale*.
4. Rudzki, T. (2026, February 12). *ChatGPT searches in English, even when you don't*. Peec AI.
5. Chirkova, N., Rau, D., Déjean, H., Formal, T., Clinchant, S., & Nikoulina, V. (2024). *Retrieval-augmented generation in multilingual settings*. arXiv.
6. Omar, M., Sorin, V., Collins, J. D., Reich, D., Freeman, R., Gavin, N., Charney, A., Stump, L., Bragazzi, N. L., Nadkarni, G. N., & Klang, E. (2025). *Multi-model assurance analysis showing large language models are highly vulnerable to adversarial hallucination attacks during clinical decision support*. *Communications Medicine*, 5, Article 330.
7. Stanford Legal Design Lab. (2025, February 11). *AI, machine translation, and access to justice*. Justice Innovation.
8. Pacheco, A. G. C., Cavalini, A., & Comarela, G. (2026). *Echoes of power: Investigating geopolitical bias in US and China large language models*. *Humanities and Social Sciences Communications*, 13, Article 675.
9. Greenberg, N. (2026, March 16). *Coverage breadth: The latest GEO research on brand visibility*. Stacker.
10. Brunswick Group. (2026, January). *2026 investor survey: Decoding investor decision-making*.
11. Chen, M., Wang, X., Chen, K., & Koudas, N. (2025). *Generative engine optimization: How to dominate AI search*. arXiv.

Research direction and framing informed by parallel findings from the University of Toronto and NAVER Labs on earned-media authority and language-dependent LLM citation patterns.^{5,11}