



HEZRON INSIGHTS

RESEARCH REPORT | AUGUST 2026

2026

African Humanitarian AI Citation Index

Whose evidence is visible in African
humanitarian crises?

Batch 1 report based on 90 production questions across four generative AI
platforms.

Hezron Ochiel | Hezron Insights, Nairobi, Kenya

5.5%

COUNTRY-LOCAL CITATIONS

hezroninsights.com

Generative AI is increasingly curating the visible authority behind humanitarian answers.

It is curating a visible hierarchy of humanitarian authority.

5.5%

of citations came from country-local producers

68.8%

came from UN or multilateral producers

40.3%

of all citation attention went to the top five producers

How the audit was built

A standardised observational audit of visible source selection in generative AI answers.

1,000

question bank

100

Batch 1 questions

10

pilot questions

90

production questions

4

AI platforms

360

platform-question
observations

10 countries

- Burkina Faso
- DRC
- Kenya
- Nigeria
- South Sudan
- Chad
- Ethiopia
- Mozambique
- Somalia
- Sudan

10 humanitarian themes × 10 question types

Conflict and protection; displacement; food insecurity; health; WASH; climate hazards; access; women and children; livelihoods/shelter/basic services; response/funding/localisation.

Question types ranged from current situation and drivers to hotspots, unmet needs, response, barriers, evidence of change and outlook.

Humanitarian reform has spent a decade asking who gets resources, leadership and recognition.

Generative AI introduces another layer of authority:
discoverability.

25%

Grand Bargain target for funding to local and national responders as directly as possible.

<10%

estimated share reaching local and national actors directly or indirectly in 2024.

3.8%

estimated share reaching them directly in 2024.


Citation localisation

The degree to which an AI-generated answer about a humanitarian situation visibly attributes supporting evidence to knowledge producers based in the affected country.



The visible evidence layer is overwhelmingly international

Publisher-type composition of 1,083 recoverable answer-linked citations.




Local evidence is visible. It is rarely dominant.




Country-local producers

60 citations across 35 publishers



African-based producers

71 citations across 43 publishers



UN / multilateral producers


745 citations

Citation attention is concentrated in a small producer core

183 distinct producers were identified; citation attention remained highly uneven.

40.3%

TOP 5



58.8%

TOP 10



71.3%

TOP 20



≈22

effective number of equally weighted publishers implied by the HHI — despite 183 distinct producers appearing in the corpus.

Five organisations captured 40.3% of citation attention.


436

citations from top five



Local visibility is scarce — and concentrated

The 60 country-local citations were distributed across 35 publishers.



Most visible local producers

● Kenya Red Cross Society	11
● The Star Kenya	5
● Nigeria NEMA	3
● Ethiopian Human Rights Commission	3


Most local producers appeared only once.

Kenya breaks through. Several countries barely do.

Complete ChatGPT sensitivity analysis: country-local share of citations by country.

29.4%

of ChatGPT citations about Kenya were country-local in the complete 90-question archive.



In the complete ChatGPT archive, three countries had no country-local citation.

0%

DRC

51 citations

0%

Somalia

48 citations


0%

Sudan

30 citations

This does not show that credible local sources were unavailable. It shows they did not become visible in this particular citation chain.

What you ask about changes whose evidence appears



In the pooled recoverable cross-platform set, WASH and humanitarian response/funding/localisation each produced 13 country-local citations, while food insecurity and malnutrition produced none.

Question framing also changes local visibility.

15.9%

local share for geographic-hotspot questions
in ChatGPT.

0%

local share for unmet-needs questions in
ChatGPT.

One plausible interpretation in the paper: geographically specific questions may create more openings for national reporting, while some humanitarian themes are strongly organised around international classification and coordination products.

Platforms expose very different citation behaviour.

CHATGPT = COMPLETE
90 / 90 source capture

Uneven source recovery prevents clean head-to-head performance rankings.

Platform	Resolved obs.	Cited obs.	Zero-URL	Citations	UN share	Local share
Gemini	19	19	0	71	78.9%	1.4%
DeepSeek	80	11	69	52	34.6%	17.3%
ChatGPT	90	90	0	567	70.9%	7.2%
Perplexity	66	66	0	393	68.4%	2.3%

The complete ChatGPT archive reproduces the same substantive pattern.

7.2%

country-local citations

95% cluster-bootstrap interval: 4.2–10.7%

70.9%

UN / multilateral citations

95% interval: 65.1–76.4%

23.3%

of questions contained at least one local citation

95% interval: 14.4–32.2%

567 citations across all 90 production questions — the only platform with complete production source-level capture.

The visible URL can differ from the organisation that produced the evidence.



217

answer-linked citations were hosted on ReliefWeb.

20.0%

of the entire recoverable corpus

93.9%

of all intermediary-hosted citations

When local evidence appears, it usually supplements international authority.

84.8%

of local-citing observations also included at least one UN or multilateral source.



Dominant pattern: supplementation. Only one cited observation relied exclusively on country-local sources.

Adding local evidence changes the balance of the stack.

Local-citing observations

47.6% **6.52**

mean UN share

mean citations

Observations without a local source

73.2% **5.67**

mean UN share

mean citations

The paper describes this as diversification plus partial substitution: local-citing answers are slightly larger and materially less UN-dominated.

Four layers reveal how visible authority is distributed.

1

REPRESENTATION

Who is visibly cited?

Country-local and African-based share.

2

CONCENTRATION

How narrow is attention?

Top-5 / top-10 / top-20 and HHI.

3

INTERMEDIATION

Who hosts vs who produced?

Separate repository or mirror from original producer.

4

SOURCE-STACK SUBSTITUTION

Does local evidence change the mix?

Supplementation, diversification or replacement.

Digital visibility is becoming part of humanitarian knowledge power.

The practical aim is to make credible local evidence easier to discover, attribute and reuse without losing provenance.

01

Publish through durable URLs

Avoid evidence that disappears inside temporary links or inaccessible drives.

02

Use accessible HTML

Make key findings readable outside image-only PDFs and social posts.

03

Show dates and authorship

Clear provenance helps both people and machines identify the source.

04

Build consistent institutional identity

Stable naming, domains and metadata reduce ambiguity.

05

Use trusted repositories wisely

Syndication can expand reach and preserve original producer metadata.

Syndication can expand local authority when producer identity remains visible.

Repositories & aggregators

- Preserve original authorship in metadata.
- Keep the producing institution visible.
- Link back to the canonical source when possible.
- Distinguish hosting from knowledge production.

International organisations

- Credit national and local partners clearly.
- Preserve data provenance in joint products.
- Avoid making a local contribution invisible in a global brand wrapper.
- Support local publishing and indexing capacity.

A strong visual finding still needs disciplined interpretation.

NO BIAS CLAIM

5.5% local share is not a formal under-citation rate without a denominator of eligible sources.

UNEQUAL SOURCE RECOVERY

105 of 360 production observations lacked recoverable source-level status.

ONE ANALYST

Source classification was not independently double-coded in this wave.

QUALITY NOT YET TESTED

The study does not evaluate claim support or source quality across the full corpus.

VISIBLE CITATIONS ONLY

A model may retrieve or use a source without exposing a URL.

ENGLISH-ONLY QUERIES

The audit may advantage organisations with English-language publishing infrastructure.

ONE-DAY SNAPSHOT

Consumer interfaces and backend retrieval behaviour can change rapidly.

The next question: How can relevant local evidence compete for visibility?

The next phase tests visibility on evidentiary merit across languages, platforms and time.

-  **MULTILINGUAL REPLICATION** Test French, Portuguese, Arabic and other language environments.
-  **DOUBLE-CODING** Independently code a random 10–20% subset and report agreement.
-  **QUALITY + SUPPORT** Assess provenance, claim support, timeliness and source quality alongside localisation.
-  **LONGITUDINAL TRACKING** Repeat the audit to distinguish stable patterns from one-day platform behaviour.
-  **CANDIDATE-POOL TESTS** Build query-specific pools of eligible local and international sources to test selection conditional on availability.

Visibility is becoming a layer of humanitarian power.

The sources AI systems cite help determine which institutions appear credible enough to investigate, repeat and trust.

5.5%

country-local

68.8%

UN / multilateral

40.3%

top-five concentration

84.8%

local-citing answers also included UN / multilateral

The policy goal is a plural evidence ecosystem in which strong local sources can become discoverable, attributable and visible.

Read the pooled figures correctly

The cross-platform citation corpus is recoverable and citation-weighted, with uneven platform contributions.



105 source-capture gaps are archival gaps, not zero-citation answers. Missingness is platform-clustered, especially in Gemini and Perplexity.

References

- Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K., & Deshpande, A. (2024). GEO: Generative engine optimization. In Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (pp. 5–16). ACM. <https://doi.org/10.1145/3637528.3671900>
- Barter, D., & Sumlut, G. M. (2023). The “conflict paradox”: Humanitarian access, localisation, and (dis)empowerment in Myanmar, Somalia, and Somaliland. *Disasters*, 47(4), 849–869. <https://doi.org/10.1111/disa.12573>
- Birhane, A. (2020). Algorithmic colonisation of Africa. *SCRIPTed*, 17(2), 389–409. <https://doi.org/10.2966/scrip.170220.389>
- Birhane, A. (2021). Algorithmic injustice: A relational ethics approach. *Patterns*, 2(2), Article 100205. <https://doi.org/10.1016/j.patter.2021.100205>
- Couldry, N., & Mejias, U. A. (2019). *The costs of connection: How data is colonising human life and appropriating it for capitalism*. Stanford University Press. <https://doi.org/10.1515/9781503609754>
- Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198237907.001.0001>
- Inter-Agency Standing Committee. (2016). *The Grand Bargain: A shared commitment to better serve people in need*.
- Kakimov, R., Tan, X., Gillham, J., & Bejtíc, N. (2026). Auditing citation behaviour in AI-generated search summaries: A framework and a case study of Google AI Overviews. In L. Bouzar-Benlabiod & C. Leung (Eds.), *Proceedings of the 39th Canadian Conference on Artificial Intelligence* (Vol. 318, pp. 224–235). PMLR.
- Mohamed, S., Png, M.-T., & Isaac, W. (2020). Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33, 659–684. <https://doi.org/10.1007/s13347-020-00405-8>

References (continued)

Mugaanyi, J., Cai, L., Cheng, S., Lu, C., & Huang, J. (2024). Evaluation of large language model performance and reliability for citations and references in scholarly writing: Cross-disciplinary study. *Journal of Medical Internet Research*, 26, e52935. <https://doi.org/10.2196/52935>

Pearson, M., & Rieger, N. (2026). The state of international humanitarian funding to local and national actors. ODI Global.

Robinson, A. (2026). Localisation without transformation: Dynamics of “localisation” in South Sudan’s humanitarian arena. *Development in Practice*. Advance online publication. <https://doi.org/10.1080/09614524.2026.2662955>

Sünkler, S., Lewandowski, D., Schultheiß, S., & Koop, O. (2026). Information diversity and authority shifts in Google’s AI Overviews: An audit of health queries. In *Companion Publication of the 2026 18th ACM Web Science Conference* (pp. 24–28). ACM. <https://doi.org/10.1145/3795513.3807429>

Venkit, P. N., Laban, P., Zhou, Y., Mao, Y., & Wu, C.-S. (2024). Search engines in an AI era: The false promise of factual and verifiable source-cited responses [Preprint]. arXiv. <https://doi.org/10.48550/arXiv.2410.22349>

Walters, W. H., & Wilder, E. I. (2023). Fabrication and errors in the bibliographic citations generated by ChatGPT. *Scientific Reports*, 13, Article 14045. <https://doi.org/10.1038/s41598-023-41032-5>

Willitts-King, B., Majid, N., Ali, M., & Poole, L. (2018). Funding to local humanitarian actors: Evidence from Somalia and South Sudan. Overseas Development Institute.

Xu, H., Iqbal, U., & Montgomery, J. M. (2026). Measuring Google AI Overviews: Activation, source quality, claim fidelity, and publisher impact [Preprint]. arXiv. <https://doi.org/10.48550/arXiv.2605.14021>



HEZRON INSIGHTS
RESEARCH | VISIBILITY | AUTHORITY

About the research

The African Humanitarian AI Citation Index (AHACI) is a structured study of source selection across humanitarian questions involving ten African countries. Batch 1 tests whose evidence becomes visibly attributable when users ask generative AI systems to explain crises.

Author
Hezron Ochiel
Hezron Insights, Nairobi, Kenya

Contact
info@hezroninsights.com
www.hezroninsights.com

Whose evidence becomes visible?

Citation localisation makes that question measurable.