Citation Evidence Report

EB-2 NIW Petition — National Interest Waiver

Matter of Dhanasar · Prong 2 (well-positioned)

MAYANK JINDAL

Machine learning engineer/Software Engineer

Google Scholar profile

Generated 2026-05-21 by CiteMap. This report organises Google Scholar citation data into the structure USCIS adjudicators apply to Prong 2 of Matter of Dhanasar (the petitioner is well positioned to advance the proposed endeavor) — the prong where past citation evidence is most probative. It is a drafting aid for the petitioner's counsel — not legal advice, and not a guarantee of any outcome. All figures must be verified, and citation counts re-snapshotted as of the petition filing date, before use in a filing.

A. Overview & Filtering Statement

45	45	7	2
Citing papers mapped	Citation edges	Home papers mapped	h-index (GS)

Filtering statement — methodology & limits

Citation **independence** is classified per citing paper by comparing the citing paper's authors to this scholar. *Self* citations are those where the scholar is an author of the citing work; *co-author* citations are by the scholar's known collaborators; *same-institution* citations are by authors affiliated with the scholar's institution(s); all remaining classified citations are *independent*. Per AAO practice, only independent citations are treated as probative of influence beyond the scholar's own circle.

Known limitations — counsel must verify. (1) Collaborator identification draws on the co-author list published on the Google Scholar profile; a collaborator not listed there may be missed, so the independent share below should be read as an **upper bound**. (2) Citation counts are a crawl-time snapshot; eligibility is judged as of the petition filing date and post-filing citations carry no weight — re-snapshot before filing. (3) Citations that could not be classified (no author data) are excluded from the percentages and reported separately.

B. Citation Independence

The AAO credits citations only where they show influence **beyond the scholar's own circle**. Self-citations and co-author citations are expressly discounted; the independent share below is the load-bearing figure.

100.0% independent of 24 classified citing papers

Citation type	Count
Independent	24
Self-citation	0
Co-author	0
Same-institution	0

21 citing papers could not be classified (no author data) and are excluded from the percentages above.

C. Significant Contributions & Their Citation Evidence

Each contribution below is presented as the AAO expects: a specific claim, followed by the **independent** citation evidence for the paper(s) that carry it. Citation counts are stated **per article**, never as a body-of-work total — the AAO holds aggregate totals to be a final-merits signal, not Criterion-5 evidence.

Where the data allows, a paper also shows its **field-normalised** standing — how its citation count ranks against Semantic Scholar papers in the same field and publication year. The comparison field is named explicitly; counsel should confirm it is the appropriate one, as the AAO scrutinises a petitioner's choice of comparison field.

Contribution 1

Claim — Contribution 1

The researcher critically examined whether large language models possess genuine causal understanding, establishing a foundational framework for evaluating causality in AI systems.

The researcher's contribution centers on the 2024 paper 'Cause and effect: Can large language models truly understand causality?', which serves as the core of this line of work. This study addresses the critical question of whether current AI architectures can genuinely comprehend causal relationships rather than merely correlating data.

This work appears to fill a significant gap in AI interpretability by challenging assumptions about LLM capabilities. By focusing on the distinction between correlation and causation, the research provides a necessary theoretical lens for assessing the limits of generative models.

The significance of this contribution is evidenced by its rapid uptake, with 58 citations recorded. Notably, 100% of the classified citing papers originate from independent researchers, indicating that the work has sparked broad, external interest and validation within the scientific community.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 23

CORE PAPER

Cause and effect: Can large language models truly understand causality?

2024 · Proceedings of the AAAI Symposium Series 4 (1), 2-9, 2024 · 58 citations (GS)

Field-normalised: 47 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2024 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	A Review of <scp>TRiSM</scp> Frameworks in Artificial Intelligence Systems: Fundamentals, Taxonomy, Use Cases, Key Challenges and Future Directions	Sikkim University	India	—
2	Biomedical natural language processing in the era of large language models	Microsoft Research	United States	—
3	Alcm: Autonomous llm-augmented causal discovery framework	University of California Irvine, University of California, Irvine	United States	—
4	ExpliCa: Evaluating explicit causal reasoning in large language models	The Hong Kong Polytechnic University, University of Pisa	Hong Kong, Italy	—
5	Counterbench: A benchmark for counterfactuals reasoning in large language models	Case Western Reserve University, Rutgers University	United States	—
6	Causal sufficiency and necessity improves chain-of-thought reasoning	City University of Hong Kong, Peking University, The Hong Kong University of Science and Technology (Guangzhou)	China, United Kingdom	—
7	C2P: Featuring large language models with causal reasoning	Carnegie Mellon University, Harvard Medical School, Sharif University of Technology	Iran, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
8	Event causality identification with synthetic control	Allen Institute for Artificial Intelligence, University of Pennsylvania	United States	—
9	GraphLocator: Graph-guided Causal Reasoning for Issue Localization	ByteDance, Peking University	China	—
10	Reasoning elicitation in language models via counterfactual feedback	Cornell Tech, Harvard University, Microsoft Research	United States	—
11	Multimodal causal reasoning benchmark: Challenging vision large language models to discern causal links across modalities	The University of Sydney	Australia	—
12	Causal-aware Large Language Models: Enhancing Decision-Making Through Learning, Adapting and Acting	Guangdong University of Technology, Huawei	China, France	—
13	CARE: Turning LLMs Into Causal Reasoning Expert	Duke University	United States	—
14	Multimodal Causal Reasoning Benchmark: Challenging Multimodal Large Language Models to Discern Causal Links Across Modalities	The University of Sydney	Australia	—
15	Causal Distillation: Transferring Structured Explanations from Large to Compact Language Models	Stanford University	United States	—
16	GraphRAG-Causal: A novel graph-augmented framework for causal reasoning and annotation in news	FAST School of Computing, NUCES Islamabad, Middle East College, University of Essex	Oman, Pakistan, United Kingdom	—
17	HCR-Reasoner: Synergizing Large Language Models and Theory for Human-like Causal Reasoning	Peking University, University of Electronic Science and Technology of China	China	—
18	Assessing LLM Reasoning Through Implicit Causal Chain Discovery in Climate Discourse	KU Leuven, Università della Svizzera italiana	Belgium, Switzerland	—
19	Causal Understanding by LLMs: The Role of Uncertainty	SUPSI	Switzerland	—
20	Benchmarking LLMs for Pairwise Causal Discovery in Biomedical and Multi-Domain Contexts	Indiana University	United States	—
21	Estimating causal effects of text interventions leveraging LLMs	University of Southern California	United States	Background
22	MUSE: A Multimodal, Generative, and Symbolic Framework for Human Experience Modeling	University of Nebraska-Lincoln	United States	—
23	CounterBench: Evaluating and Improving Counterfactual Reasoning in Large Language Models	Case Western Reserve University, Rutgers University	United States	—

Independent citing papers only; self- and co-author citations excluded. The **S2** column carries Semantic Scholar's read of each citation — *Methodology / Result* (the citing work used the method or built on the finding — the "built on / relied upon" pattern the AAO credits), *Influential* (S2's isInfluential signal, Valenzuela et al. 2015), or *Background* (a passing mention).

D. Citing-Institution Prestige & Geography

Top citing institutions

Institution	Country	World ranking	Citing papers
Peking University	China	SCImago #11 · THE 13 · QS 14	3
Rutgers University	United States	—	2
Case Western Reserve University	United States	SCImago #627 · THE =145 · QS =294	2
University of California, Irvine	United States	SCImago #329 · THE 97 · QS 293	2
The University of Sydney	Australia	SCImago #93 · THE =53 · QS =25	2
Microsoft Research	United States	—	2
Sharif University of Technology	Iran	SCImago #4501 · THE 351–400 · QS =375	1
Cornell Tech	United States	—	1
Tianjin University	China	SCImago #90 · THE 201–250 · QS =257	1
The Hong Kong Polytechnic University	Hong Kong	SCImago #256 · THE 80 · QS 54	1
The Hong Kong University of Science and Technology (Guangzhou)	China	SCImago #483 · THE =58 · QS 44	1
KU Leuven	Belgium	SCImago #180 · THE 46 · QS 60	1
University College London	United Kingdom	SCImago #30	1
University of California Irvine	United States	SCImago #329 · THE 97 · QS 293	1
Harvard University	United States	SCImago #4 · THE =5 · QS 5	1

Geographic distribution of citing authors

Country	Citing papers
United States	13
China	4
Australia	2
Switzerland	2
United Kingdom	2
India	1
Iran	1
Italy	1
Oman	1
Pakistan	1
Belgium	1
France	1

Citing-institution prestige and the spread of citing countries speak to recognition **beyond the scholar's own institution and circle** — the dispersion the AAO looks for. World rankings (SCImago / THE / QS) are context, not a stand-alone criterion: the AAO does not treat a citing institution's rank as probative on its own.

F. AAO Precedent Considerations

Pre-filing self-check (AAO denial patterns)

The AAO non-precedent decisions reject citation evidence on a small set of recurring grounds. Confirm the petition addresses each before filing:

- ☐ Self-citations are disclosed and netted out — a Google Scholar total alone is faulted (§1.1).
- ☐ Evidence is per individual article, not a body-of-work aggregate total (§1.2).
- ☐ The petition articulates why the citations show major significance — numbers never stand alone (§1.5).
- ☐ For the strongest papers, citation content shows the work was built on / relied upon, not just listed (§1.6, §2.2).
- ☐ Co-author / collaborator citations are identified and not counted as independent (§1.7).
- ☐ Recognition is shown beyond the scholar's own institution and circle (§1.8).
- ☐ Every citation figure is snapshotted as of the filing date; post-filing citations are excluded (§1.9).
- ☐ Journal impact factor / downloads are not relied on as proxies for article significance (§1.10, §1.12).
- ☐ For large-collaboration papers, the scholar's specific role is documented (§1.13).
- ☐ Aggregate totals / h-index / field-relative rates are placed in a clearly-labelled final-merits section, per Kazarian (§3, §6.1.7).

Disclaimer

The AAO decisions referenced here are **non-precedent** — persuasive illustrations of how USCIS reasons, not binding law. This report is a drafting aid produced from public citation data; it is not legal advice and does not assess the petition's merits. All analysis must be reviewed by qualified immigration counsel.

G. Citation Evidence Index

Cross-reference of each contribution to the regulatory criterion it supports. Counsel should map these to the petition's exhibit numbers.

Contribution	Core paper	Indep. cites	Supports
Contribution 1	Cause and effect: Can large language models truly understand causality?	23	Dhanasar — Prong 2 (well-positioned)