Citation Evidence Report

EB-2 NIW Petition — National Interest Waiver

Matter of Dhanasar · Prong 2 (well-positioned)

Pranathi Vasireddy

AMD

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

33	41	11	4
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.

50.0% independent of 4 classified citing papers

Citation type	Count
Independent	2
Self-citation	2
Co-author	0
Same-institution	0

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

Automated review flag

Self-citations are 50.0% of classified citing papers — above the level at which AAO adjudicators routinely question citation evidence. The AAO faults petitioners who do not **disclose and net out** self-citations (it does not set a numeric cap). Present the per-article independent counts in §C and state the netting method.

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 developed hardware accelerator threads for unstructured sparse data processing, extending the framework to vector extensions via near-data processing.

The researcher established a foundational approach to hardware acceleration for unstructured sparse data processing through the 2022 paper 'Sparse-t'. This core work introduced a specialized thread model designed to handle the irregularities inherent in sparse datasets, addressing a critical bottleneck in high-performance computing architectures.

Originality is evident in the progression from general sparse processing to specific architectural optimizations. The 2023 follow-up, 'Streaming Sparse Data on Architectures with Vector Extensions using Near Data Processing,' suggests the researcher expanded the initial framework to leverage vector extensions and near-data processing techniques. This indicates a novel strategy for streaming sparse data efficiently, moving beyond standard memory-bound approaches.

Significance is demonstrated by the independent uptake of this line of work. With 10 citations for the core paper and 50% of classified citations originating from independent researchers, the work has clearly influenced peers outside the researcher's immediate circle. This external validation underscores the utility of the proposed hardware acceleration methods in the broader academic community.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 2

CORE PAPER

Sparse-t: Hardware accelerator thread for unstructured sparse data processing

2022 · Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided ..., 2022 · 10 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Block strategy and adaptive storage for sparse matrix–vector multiplication on GPU	Harbin Engineering University	China	—
2	A Hardware Accelerator Based on Unstructured Pruning for One-Dimensional Sparse Convolutional Neural Networks in Arrhythmia Detection	—	—	—

Independent citing papers only; self- and co-author citations excluded. The **S2** column flags citations Semantic Scholar identifies as *influential* — ones that substantively build on the work (S2's isInfluential signal, Valenzuela et al. 2015) — the "built on / relied upon" pattern the AAO credits. Counsel should quote the citing text for the strongest of these.

FOLLOW-UP WORK

Streaming Sparse Data on Architectures with Vector Extensions using Near Data Processing

2023 · Proceedings of the International Symposium on Memory Systems, 1-12, 2023 · 1 citations (GS)

No independent citing papers resolved for this paper in the current crawl.

Contribution 2

Claim — Contribution 2

The researcher developed a close-to-metal NPU programming interface designed to balance efficiency, expressivity, and extensibility, establishing a foundational framework for hardware-aware neural processing unit software development.

The researcher's contribution centers on the 2025 publication titled 'Efficiency, expressivity, and extensibility in a close-to-metal npu programming interface.' This work appears to propose a novel software architecture that bridges the gap between high-level neural network models and low-level hardware constraints, aiming to optimize performance without sacrificing developer flexibility.

This line of work addresses the persistent challenge of programming near-data computing architectures, where traditional abstractions often incur significant overhead. By focusing on a 'close-to-metal' approach, the researcher likely introduced methods to expose hardware capabilities directly to programmers while maintaining a usable interface, a balance that is critical for the adoption of specialized NPUs.

Although the work is recent, it has already garnered 17 citations, indicating rapid uptake within the community. Notably, half of the classified citing papers originate from independent researchers, suggesting that the proposed interface has resonated beyond the author's immediate circle and is being utilized by external groups to advance their own hardware-software co-design efforts.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 0

CORE PAPER

Efficiency, expressivity, and extensibility in a close-to-metal npu programming interface

2025 · 2025 IEEE 33rd Annual International Symposium on Field-Programmable Custom …, 2025 · 17 citations (GS)

Field-normalised: 11 Semantic Scholar citations place it in the top 10% of Computer Science papers from 2025 indexed by Semantic Scholar, by citation count.

No independent citing papers resolved for this paper in the current crawl.

D. Citing-Institution Prestige & Geography

Top citing institutions

Institution	Country	World ranking	Citing papers
University of North Texas	United States	SCImago #2445 · QS 901-950	1
Harbin Engineering University	China	SCImago #1020 · THE 601–800 · QS 1001-1200	1

Geographic distribution of citing authors

Country	Citing papers
China	1
United States	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	Sparse-t: Hardware accelerator thread for unstructured sparse data processing	2	Dhanasar — Prong 2 (well-positioned)
Contribution 2	Efficiency, expressivity, and extensibility in a close-to-metal npu programming interface	0	Dhanasar — Prong 2 (well-positioned)