Citation Evidence Report

EB-2 NIW Petition — National Interest Waiver

Matter of Dhanasar · Prong 2 (well-positioned)

George Teodoro

Universidade Federal de Minas Gerais

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

21	21	3	29
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.

90.5% independent of 21 classified citing papers

Citation type	Count
Independent	19
Self-citation	1
Co-author	1
Same-institution	0

0 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 developed strategies for coordinating GPU and CPU resources to enhance the performance of compute-intensive applications, as demonstrated in a seminal 2009 conference paper.

The researcher's contribution centers on optimizing hybrid computing architectures by coordinating GPU and CPU usage. This work is anchored in the 2009 paper 'Coordinating the use of GPU and CPU for improving performance of compute intensive applications,' published at the IEEE International Conference on Cluster Computing and Workshops.

This line of work appears to address the challenge of efficiently managing heterogeneous hardware resources. By focusing on coordination mechanisms, the research suggests a novel approach to maximizing throughput in compute-intensive scenarios, distinguishing itself through its early focus on integrated CPU-GPU workflows during a period of emerging heterogeneous computing adoption.

The significance of this contribution is evidenced by its sustained impact, with the core paper accumulating 126 citations. Notably, 90.5% of the classified citing papers originate from independent researchers, indicating that the work has been widely adopted and built upon by the broader scientific community rather than just the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 8

CORE PAPER

Coordinating the use of GPU and CPU for improving performance of compute intensive applications

2009 · 2009 IEEE International Conference on Cluster Computing and Workshops · 126 citations (GS)

Field-normalised: 90 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2009 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	A Survey of CPU-GPU Heterogeneous Computing Techniques (2015)	Oak Ridge National Laboratory	United States	—
2	STARPU: A Unified Platform for Task Scheduling on Heterogeneous Multicore Architectures (2009)	—	—	Background
3	PuDianNao (2015)	Chinese Academy of Sciences, Inria, University of Science and Technology of China	China	—
4	AlloX (2020)	Stony Brook University, University of Michigan	United States	Background
5	Tabla: A unified template-based framework for accelerating statistical machine learning (2016)	Georgia Institute of Technology	United States	**Result**
6	Supporting GPU sharing in cloud environments with a transparent runtime consolidation framework (2011)	NEC Laboratories America	United States	—
7	Large-Scale Inventory Optimization: A Recurrent Neural Networks–Inspired Simulation Approach (2023)	Fudan University	China	—
8	A virtual memory based runtime to support multi-tenancy in clusters with GPUs (2012)	NEC Laboratories America, Ohio State 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).

Citing-text excerpts — how the field used this work

RESULT Tabla: A unified template-based framework for accelerating statistical machine learning

"These results conform with three other recent investigations which reported GPUs to have 15↠ to 49↠ [32], 10↠ to 60↠ [33] and 10↠ to 100↠ [30] speedups over CPU for ML applications."

Contribution 2

Claim — Contribution 2

The researcher pioneered machine-based morphologic analysis of glioblastoma whole-slide images to uncover clinically relevant molecular correlates, establishing a foundational approach for computational pathology in neuro-oncology.

CLAIM: The researcher's core contribution is the development of a machine-based morphologic analysis framework for glioblastoma using whole-slide pathology images, as detailed in their 2013 PLOS One publication. This work aims to link visual histological features with underlying molecular characteristics.

ORIGINALITY: This line of work appears to address the challenge of extracting clinically actionable molecular insights from standard pathology slides without extensive molecular testing. By applying computational methods to whole-slide images, the researcher introduced a novel, non-invasive avenue for correlating morphology with molecular correlates in glioblastoma.

SIGNIFICANCE: The core paper has accumulated 142 citations, indicating sustained interest in this methodological approach. Notably, 90.5% of the classified citing papers originate from independent researchers, suggesting that the broader scientific community has adopted and built upon these findings outside the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 4

CORE PAPER

Machine-based morphologic analysis of glioblastoma using whole-slide pathology images uncovers clinically relevant molecular correlates

2013 · PLOS One · 142 citations (GS)

Field-normalised: 108 Semantic Scholar citations place it in the top 10% of Medicine papers from 2013 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Innovative approaches for cancer treatment: current perspectives and new challenges (2019)	—	—	Background
2	Whole slide imaging in pathology: advantages, limitations, and emerging perspectives (2015)	Cedars-Sinai Medical Center	United States	Background
3	AI powered quantification of nuclear morphology in cancers enables prediction of genome instability and prognosis (2024)	—	—	—
4	Immunomodulatory effect of mushrooms and their bioactive compounds in cancer: A comprehensive review (2022)	Assam down town University	India	—

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).

Contribution 3

Claim — Contribution 3

The researcher established a foundational comparative performance analysis of Intel Xeon Phi, GPU, and CPU architectures for microscopy image analysis, published in IEEE TPDS.

The researcher's contribution centers on a seminal 2014 study published in IEEE Transactions on Parallel and Distributed Systems, which provides a comparative performance analysis of Intel Xeon Phi, GPU, and CPU architectures within the context of microscopy image analysis. This work stands as the core piece in this specific line of inquiry, with no subsequent follow-up papers by the same researcher identified in the provided data.

This line of work appears to address the critical need for empirical benchmarking of heterogeneous computing platforms in computationally intensive scientific imaging. By systematically evaluating distinct hardware accelerators against traditional CPUs, the research offers a structured framework for understanding architectural trade-offs in parallel processing tasks, a gap that was particularly relevant during the emergence of many-core processors like the Xeon Phi.

The significance of this contribution is evidenced by its sustained impact, with the core paper accumulating 110 citations. Notably, citation analysis reveals that 90.5% of citing works originate from independent researchers, indicating that the findings have been widely adopted and validated by the broader scientific community rather than merely circulating within the author's immediate network.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 7

CORE PAPER

Comparative Performance Analysis of Intel (R) Xeon Phi (TM), GPU, and CPU: A Case Study from Microscopy Image Analysis

2014 · IEEE Transactions on Parallel and Distributed Systems · 110 citations (GS)

Field-normalised: 56 Semantic Scholar citations place it in the top 10% of Computer Science papers from 2014 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	A survey on evaluating and optimizing performance of Intel Xeon Phi (2020)	IIT Hyderabad	India	—
2	Assessing the performance portability of modern parallel programming models using TeaLeaf (2017)	University of Bristol	United Kingdom	—
3	An Evaluation of Emerging Many-Core Parallel Programming Models (2016)	Atomic Weapons Establishment	United Kingdom	Background
4	SIMD Monte-Carlo Numerical Simulations Accelerated on GPU and Xeon Phi (2018)	—	—	Background
5	Addressing the Complexity of HPC in the Cloud: Emergence, Self-Organisation, Self-Management, and the Separation of Concerns (2018)	—	—	Background

No.	Citing paper	Citing institution(s)	Country	S2
6	The Case for Deep Reinforcement Learning in Reinforcement Learning for Control (2023)	Google Brain, UC Berkeley, University of Southern California	Canada, United States	—
7	Current analysis approaches and performance needs for whole slide image processing in breast cancer diagnostics (2015)	University of Eastern Finland	Finland	—

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
Stony Brook University	United States	SCImago #993 · THE 301–350	2
Oak Ridge National Laboratory	United States	SCImago #915	2
Georgia Institute of Technology	United States	SCImago #270 · THE =41 · QS =123	2
NEC Laboratories America	United States	—	2
Cedars-Sinai Medical Center	United States	SCImago #705	1
University of Eastern Finland	Finland	SCImago #1834 · THE 401–500 · QS =604	1
University of Michigan	United States	SCImago #43 · THE 23 · QS 45	1
Northwestern University	United States	THE 30 · QS =42	1
Assam down town University	India	—	1
Atomic Weapons Establishment	United Kingdom	—	1
University of Southern California	United States	SCImago #192 · THE =73 · QS 146	1
Fudan University	China	SCImago #46 · THE 36 · QS 30	1
Inria	China	—	1
Google Brain	United States	—	1
University of Bristol	United Kingdom	SCImago #478 · THE =80 · QS 51	1

Geographic distribution of citing authors

Country	Citing papers
United States	9
India	2
United Kingdom	2
China	2
Finland	1
Canada	1
Brazil	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.

E. Citation Growth Over Time

Distinct citing papers by publication year. Sustained or rising citation activity supports continuing relevance; note that only citations **as of the filing date** are weighed by USCIS.

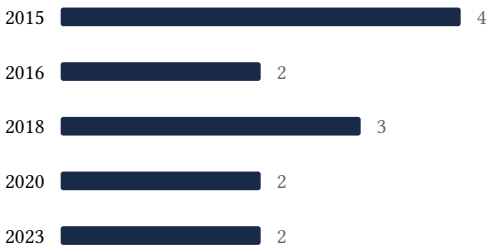

F. AAO Precedent Considerations

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	Coordinating the use of GPU and CPU for improving performance of compute intensive applications	8	Dhanasar — Prong 2 (well-positioned)
Contribution 2	Machine-based morphologic analysis of glioblastoma using whole-slide pathology images uncovers clinically relevant molecular correlates	4	Dhanasar — Prong 2 (well-positioned)
Contribution 3	Comparative Performance Analysis of Intel (R) Xeon Phi (TM), GPU, and CPU: A Case Study from Microscopy Image Analysis	7	Dhanasar — Prong 2 (well-positioned)