Citation Evidence Report

EB-1B Petition — Outstanding Professor or Researcher

8 CFR § 204.5(i)(3) · Authorship + Original Contributions

Kurt Keutzer

Professor of the Graduate School, EECS, University of California, Berkeley

Google Scholar profile

Generated 2026-06-11 by CiteMap. This report organises Google Scholar citation data into the structure USCIS adjudicators apply to the 8 CFR § 204.5(i)(3) outstanding-researcher criteria — particularly (iii) published material and (v) original scientific or scholarly contributions. 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

791	791	30	123
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.

96.8% independent of 778 classified citing papers

Citation type	Count
Independent	753
Self-citation	3
Co-author	22
Same-institution	0

13 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 pioneered lightweight neural network architectures, notably SqueezeNet, and advanced hardware-aware design and quantization methods to enable efficient, low-power deep learning inference.

The researcher's core contribution rests on the 2016 paper SqueezeNet, which achieved AlexNet-level accuracy with significantly fewer parameters and a sub-megabyte model size. This work established a foundation for efficient deep learning by demonstrating that high performance could be maintained while drastically reducing model complexity.

This line of work appears to address the critical gap between high-accuracy deep learning models and the computational constraints of real-world deployment. The subsequent titles suggest a logical progression from architectural compression to broader efficiency strategies. The 2018 paper FBNet indicates an expansion into hardware-aware design via differentiable neural architecture search, while the 2022 survey on quantization methods reflects a continued focus on optimizing inference for low-power computer vision applications.

The significance of this research is evidenced by substantial citation counts, with the core paper accumulating over 12,000 citations and follow-up works receiving nearly 2,000 and 2,500 citations respectively. Furthermore, analysis of citing papers reveals that 96.8% originate from independent researchers, indicating that this work has been widely adopted and validated by the broader scientific community rather than relying on self-citation or institutional bias.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 753 · **3 flagged influential by Semantic Scholar**

CORE PAPER

SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and< 0.5 MB model size

2016 · arXiv preprint arXiv:1602.07360, 2016 · 12,675 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Object detection in 20 years: A survey	Beihang University, Carleton University, Nanyang Technological University	Canada, China, Singapore	—
2	A comprehensive review of convolutional neural networks for defect detection in industrial applications	University of Huddersfield	United Kingdom	—
3	Deep learning for 3d point clouds: A survey	Inception Institute of Artificial Intelligence, Information Engineering University, Nanyang Technological University	Australia, China, Singapore	—
4	Deep learning for generic object detection: A survey	Arizona State University, National University of Defense Technology, Peking University	Canada, China, Finland	—
5	A comprehensive survey of loss functions and metrics in deep learning	Centro de Investigaciones en Óptica A.C., Instituto Politecnico Nacional, Universidad Autónoma de Querétaro	Mexico	—
6	Deep face recognition: A survey	Beijing University of Posts and Telecommunications	China	—
7	The elements of end-to-end deep face recognition: A survey of recent advances	Ryerson University, Shanghai University	Canada, China	—

No.	Citing paper	Citing institution(s)	Country	S2
8	Deep learning for medical image-based cancer diagnosis	Nanjing Normal University, University of Leicester	China, United Kingdom	—
9	A survey on deep learning for cybersecurity: Progress, challenges, and opportunities	Universidad de las Fuerzas Armadas ESPE, Zhejiang University	China, Ecuador	—
10	End-edge-cloud collaborative computing for deep learning: A comprehensive survey	Beijing Institute of Technology, Simula Metropolitan Center for Digital Engineering, University of Chinese Academy of Sciences	China, Norway	—
11	Neural networks and deep learning	Brown University, University of Pennsylvania	United States	—
12	A survey of machine learning and deep learning in remote sensing of geological environment: Challenges, advances, and opportunities	China University of Geosciences	China	—
13	Explaining deep neural networks and beyond: A review of methods and applications	Berlin Institute for the Foundations of Learning and Data, Fraunhofer Institute for Telecommunications, Heinrich Hertz Institute	Germany	—
14	Review of image classification algorithms based on convolutional neural networks	Guizhou University	China	—
15	Edge deep learning in computer vision and medical diagnostics: a comprehensive survey	University of New South Wales, UNSW Sydney	Australia	—
16	Pruning and quantization for deep neural network acceleration: A survey	University of Science and Technology Beijing	China	—
17	A state-of-the-art survey on deep learning theory and architectures	Comcast, Lawrence Livermore National Laboratory, Saint Louis University	United States	—
18	Efficient processing of deep neural networks: A tutorial and survey	Massachusetts Institute of Technology	United States	—
19	Convergence of edge computing and deep learning: A comprehensive survey	Huawei Technologies, Nanyang Technological University, Shenzhen University	China, Singapore	—
20	Classification of COVID-19 in chest X-ray images using DeTraC deep convolutional neural network	Assiut University, Birmingham City University	Egypt, United Kingdom	—
21	Deep learning in mobile and wireless networking: A survey	Imperial College London, Microsoft, University of Edinburgh	United Kingdom	—
22	Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges	Karlsruhe Institute of Technology, Robert Bosch GmbH, Ulm University	Germany	—
23	Deep residual learning in spiking neural networks	Centre de Recherche Cerveau et Cognition, Peking University, Peng Cheng Laboratory	China, France	—
24	On the analyses of medical images using traditional machine learning techniques and convolutional neural networks	Beijing University of Technology, Prince Sultan University, University of Central Punjab	China, Kingdom of Saudi Arabia, Pakistan	—

No.	Citing paper	Citing institution(s)	Country	S2
25	Harnessing multimodal data integration to advance precision oncology	Memorial Sloan Kettering Cancer Center	United States	—
26	Cellular, wide-area, and non-terrestrial IoT: A survey on 5G advances and the road toward 6G	Aalborg University, Ericsson (Sweden), Interdisciplinary Centre for Security, Reliability and Trust	Australia, Denmark, Luxembourg	—
27	A review of object detection based on deep learning	Xi'an Jiaotong University, Xi'an Jiaotong University	China	—
28	Edge machine learning for ai-enabled iot devices: A review	University of Reggio Calabria	Italy	—
29	A comprehensive review of model compression techniques in machine learning: PV Dantas et al.	The University of Manchester, Universidade Federal do Amazonas	Brazil, United Kingdom	—
30	A survey of deep learning techniques for autonomous driving	Transilvania University of Brasov	Romania	—

Showing the 30 most-cited of 753 independent citing papers.

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).

FOLLOW-UP WORK

A survey of quantization methods for efficient neural network inference

2022 · Low-power computer vision, 291-326, 2022 · 2,449 citations (GS)

Field-normalised: 1,496 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2022 indexed by Semantic Scholar, by citation count.

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

FOLLOW-UP WORK

FBNet: Hardware-Aware Efficient ConvNet Design via Differentiable Neural Architecture Search

2018 · arXiv, 2018 · 1,968 citations (GS)

Field-normalised: 1,422 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2018 indexed by Semantic Scholar, by citation count.

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

Contribution 2

Claim — Contribution 2

The researcher developed DAGON, a seminal 1987 framework for technology binding and local optimization via DAG matching, establishing a foundational approach in electronic design automation.

CLAIM: The researcher's primary contribution is the development of DAGON, introduced in a 1987 paper published in the Proceedings of the 24th ACM/IEEE Design Automation Conference. This work presents a method for technology binding and local optimization through directed acyclic graph (DAG) matching.

ORIGINALITY: The title suggests this work addressed the technical challenge of mapping logical designs to specific technology libraries efficiently. By focusing on DAG matching, the researcher appears to have introduced a structured algorithmic approach to local optimization, distinguishing it from broader or less precise binding techniques prevalent at the time.

SIGNIFICANCE: The core paper has accumulated 802 citations, indicating substantial influence in the field. Notably, 96.8% of the citing papers originate from independent researchers, demonstrating that the methodology was widely adopted and validated by the broader scientific community rather than just the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 0

CORE PAPER

DAGON: Technology binding and local optimization by DAG matching

1987 · Proceedings of the 24th ACM/IEEE Design Automation Conference, 341-347, 1987 · 802 citations (GS)

Field-normalised: 327 Semantic Scholar citations place it in the top 5% of Computer Science papers from 1987 indexed by Semantic Scholar, by citation count.

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

Contribution 3

Claim — Contribution 3

The researcher provided a seminal, highly cited synthesis of the parallel computing research landscape, establishing a foundational reference point for the field.

CLAIM: The researcher's contribution centers on the 2006 technical report titled 'The landscape of parallel computing research: A view from berkeley,' which serves as the core work in this line of inquiry. This document appears to offer a comprehensive overview or synthesis of the state of parallel computing at that time.

ORIGINALITY: Given the title and the nature of technical reports from major institutions, this work likely addressed a need for a consolidated view of a fragmented or rapidly evolving research area. By framing the research 'from Berkeley,' the researcher appears to have provided a distinctive perspective or authoritative summary that distinguished itself from scattered individual studies, offering a structured understanding of the field's landscape.

SIGNIFICANCE: The work has achieved substantial impact, evidenced by 3,111 citations. Notably, analysis of 778 citing papers reveals that 96.8% originate from independent researchers, indicating that the contribution was widely adopted and valued by the broader scientific community rather than just the researcher's immediate circle. This high degree of independent citation suggests the work serves as a standard reference or foundational text in parallel computing.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 0

CORE PAPER

The landscape of parallel computing research: A view from berkeley

2006 · Technical Report UCB/EECS-2006-183, EECS Department, University of ..., 2006 · 3,111 citations (GS)

Field-normalised: 2,459 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2006 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
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	34
University of California, Irvine Medical Center	United States	—	25
UC Berkeley	United States	—	17
Massachusetts Institute of Technology	United States	SCImago #41 · THE 2 · QS 1	16
Stanford University	United States	SCImago #18 · THE =5 · QS 3	15
Alibaba Group	United States	SCImago #226	13
Zhejiang University	China	SCImago #6 · THE 39 · QS 49	12
University of Electronic Science and Technology of China	China	SCImago #129 · THE 301–350 · QS =519	12
Nanyang Technological University	Singapore	SCImago #137	12
Chinese Academy of Sciences	China	SCImago #2	12
Peking University	China	SCImago #11 · THE 13 · QS 14	11
Google	United States	—	11
Microsoft Research	United States	—	10
Tianjin University	P. R. China	SCImago #90 · THE 201–250 · QS =257	9
Xidian University	China	SCImago #269 · THE 601–800	9

Geographic distribution of citing authors

Country	Citing papers
China	276
United States	247
United Kingdom	64
India	63
South Korea	46
Canada	41
Australia	31
Germany	28
Italy	26
Singapore	24
Turkey	24
Saudi Arabia	21

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	SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and< 0.5 MB model size	753	8 CFR 204.5(i)(3) — Outstanding Researcher
Contribution 2	DAGON: Technology binding and local optimization by DAG matching	0	8 CFR 204.5(i)(3) — Outstanding Researcher
Contribution 3	The landscape of parallel computing research: A view from berkeley	0	8 CFR 204.5(i)(3) — Outstanding Researcher