

Citation Evidence Report

EB-1B Petition — Outstanding Professor or Researcher

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


Andrew Zisserman

University of Oxford

[Google Scholar profile](#)

Generated 2026-05-21 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



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.

97.6% independent of 42 classified citing papers

Citation type	Count
Independent	41
Self-citation	0
Co-author	1
Same-institution	0

6 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 advanced large-scale image recognition by developing very deep convolutional networks, a foundational contribution evidenced by over 160,000 citations and near-universal independent adoption.

The researcher's primary contribution lies in the development of very deep convolutional networks for large-scale image recognition, as detailed in their seminal 2014 paper. This work stands as a singular, high-impact achievement in the field, with no subsequent follow-up papers by the same author required to establish its foundational nature.

This line of work appears to address the challenge of scaling neural network depth for complex visual tasks. The title suggests a methodological innovation in architecture design, enabling the effective training of deeper models than previously standard, thereby pushing the boundaries of what convolutional networks could achieve in recognition accuracy.

The significance of this contribution is underscored by its extraordinary citation count of 161,256, indicating widespread influence. Furthermore, analysis of citing literature reveals that 97.6% of citations originate from independent researchers, demonstrating that the work has been broadly adopted and validated by the global scientific community rather than relying on self-citation or institutional bias.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 10 · 2 flagged influential by Semantic Scholar

CORE PAPER

[Very deep convolutional networks for large-scale image recognition](#)

2014 · arXiv preprint arXiv:1409.1556, 2014 · 161,256 citations (GS)

Field-normalised: 110,772 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2014 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Poly Kernel Inception Network for Remote Sensing Detection	Communication University of China, Nanjing University of Science and Technology, Zhejiang University	China	—
2	Rewrite the Stars (2024)	Microsoft, Northeastern University	United States	—
3	EMCAD: Efficient Multi-scale Convolutional Attention Decoding for Medical Image Segmentation (2024)	The University of Texas at Austin	United States	—
4	YOLOv6: A Single-Stage Object Detection Framework for Industrial Applications	Meituan Inc.	—	Influential
5	Vision Mamba: Efficient Visual Representation Learning with Bidirectional State Space Model (2024)	Beijing Academy of Artificial Intelligence, Horizon Robotics, Huazhong University of Science and Technology	China	—
6	Qwen2.5-Math Technical Report: Toward Mathematical Expert Model via Self-Improvement (2024)	Alibaba Group	China	—
7	Vision-Language Models for Vision Tasks: A Survey	Nanyang Technological University, Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences	China, Singapore	Influential

No.	Citing paper	Citing institution(s)	Country	S2
8	Wavelet Convolutions for Large Receptive Fields (2024)	Ben-Gurion University of the Negev	Israel	—
9	YOLO-v1 to YOLO-v8, the Rise of YOLO and Its Complementary Nature toward Digital Manufacturing and Industrial Defect Detection	University of Huddersfield	United Kingdom	—
10	Hardware implementation of memristor-based artificial neural networks	Hebei University, IBM Research, IBM Research - Zurich	China, Italy, Saudi Arabia	—

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.

Contribution 2

Claim – Contribution 2

The researcher established a foundational theoretical framework for multiple view geometry in computer vision, codified in a seminal textbook that has become a standard reference in the field.

The researcher’s primary contribution is the systematic formulation of multiple view geometry for computer vision, anchored by the publication of the textbook 'Multiple View Geometry in Computer Vision' in 2003. This work serves as the core intellectual output for this line of research, with no subsequent follow-up papers by the same author listed in the provided data.

This contribution appears to address the need for a rigorous, unified mathematical treatment of geometric relationships across multiple camera views. By consolidating these principles into a comprehensive textbook, the researcher provided a structured foundation that likely filled a gap in the literature for both theoretical clarity and practical application in the field.

The significance of this work is evidenced by its extensive citation record, with over 37,000 citations indicating widespread adoption. Furthermore, analysis of citing papers reveals that 97.6% originate from independent researchers, suggesting that the work has had a broad, field-wide impact beyond the researcher’s immediate academic circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 5 · 1 flagged influential by Semantic Scholar

CORE PAPER

[Multiple View Geometry in Computer Vision](#)

2003 · Cambridge University Press (Textbook) · 37,540 citations (GS)

Field-normalised: 18,513 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2003 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	DUSt3R: Geometric 3D Vision Made Easy (2024)	Aalto University, Naver, Naver Labs Europe	Finland, France	—
2	LightGlue: Local Feature Matching at Light Speed	ETH Zurich, Microsoft	Switzerland	—
3	Grounding Image Matching in 3D with MAST3R (2024)	Naver, Naver Labs Europe	France	Influential
4	BLINK: Multimodal Large Language Models Can See but Not Perceive (2024)	Allen Institute for AI, Columbia University, University of California, Davis	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
5	Continuous 3D Perception Model with Persistent State	UC Berkeley	United States	—

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.

Contribution 3

Claim – Contribution 3

The researcher established the PASCAL VOC Challenge, a seminal benchmark that standardized visual object class recognition and catalyzed widespread adoption of rigorous evaluation in computer vision.

CLAIM: The researcher’s primary contribution is the establishment of the PASCAL Visual Object Classes (VOC) Challenge, published in the International Journal of Computer Vision in 2010. This work serves as the foundational core of this line of research, with no subsequent follow-up papers by the same researcher identified in the provided data.

ORIGINALITY: The title suggests this work addressed a critical need for standardized evaluation in visual object recognition. By defining a specific challenge framework, the researcher appears to have created a common ground for comparing algorithms, moving the field beyond isolated, non-comparable experiments toward a unified benchmarking approach.

SIGNIFICANCE: The work has achieved substantial impact, evidenced by over 28,000 citations. Analysis of citing literature reveals that 97.6% of classified citations originate from independent researchers, indicating that the PASCAL VOC Challenge has been widely adopted and utilized by the broader scientific community rather than just the researcher’s immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 10 · 1 flagged influential by Semantic Scholar

CORE PAPER

[The PASCAL Visual Object Classes \(VOC\) Challenge](#)

2010 · International Journal of Computer Vision (IJCV) · 28,240 citations (GS)

Field-normalised: 15,980 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2010 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	InternImage: Exploring Large-Scale Vision Foundation Models with Deformable Convolutions (2023)	Nanjing University, SenseTime, SenseTime Research	China, Hong Kong	—
2	DINOv2: Learning Robust Visual Features without Supervision	Inria, Meta	France	—
3	Parameter-Efficient Fine-Tuning for Large Models: A Comprehensive Survey (2024)	Arizona State University, City University of New York (Baruch College & Graduate Center), New York University	Canada, United States	—
4	SigLIP 2: Multilingual Vision-Language Encoders with Improved Semantic Understanding, Localization, and Dense Features	Google DeepMind	United Kingdom	—
5	Vision-Language Models for Vision Tasks: A Survey	Nanyang Technological University, Shenzhen Institutes	China, Singapore	—

No.	Citing paper	Citing institution(s)	Country	S2
		of Advanced Technology, Chinese Academy of Sciences		
6	A survey on deep neural network pruning: Taxonomy, comparison, analysis, and recommendations (2024)	Harbin Institute of Technology, Harbin Institute of Technology (Shenzhen), The University of Adelaide	Australia, China	—
7	A Comprehensive Survey on Pretrained Foundation Models: A History from BERT to ChatGPT	Beihang University, Duke University, Hangzhou Dianzi University	Australia, China, Singapore	—
8	Thinking in Space: How Multimodal Large Language Models See, Remember, and Recall Spaces	New York University, Stanford University, Yale University	United States	—
9	YOLO-v1 to YOLO-v8, the Rise of YOLO and Its Complementary Nature toward Digital Manufacturing and Industrial Defect Detection	University of Huddersfield	United Kingdom	—
10	A Comprehensive Review of YOLO Architectures in Computer Vision: From YOLOv1 to YOLOv8 and YOLO-NAS	Instituto Politecnico Nacional, Instituto Politécnico Nacional, Universidad Autónoma de Querétaro	Mexico	Influential

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.

D. Citing-Institution Prestige & Geography

Top citing institutions

Institution	Country	World ranking	Citing papers
Shanghai AI Laboratory	China	—	4
Nanjing University	China	SCImago #178 · THE =62 · QS =103	3
University of California, Berkeley	United States	SCImago #95 · THE 9 · QS =17	3
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	3
New York University	United States	SCImago #116 · THE =31 · QS 55	2
Peking University	China	SCImago #11 · THE 13 · QS 14	2
Stanford University	United States	SCImago #18 · THE =5 · QS 3	2
Naver	France	—	2
Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences	China	SCImago #403	2
University of Southern California	United States	SCImago #192 · THE =73 · QS 146	2
Meta AI	United States	—	2
Naver Labs Europe	France	—	2
Carnegie Mellon University	United States	SCImago #266 · THE 24 · QS 52	2

Institution	Country	World ranking	Citing papers
SenseTime	China	—	2
Nanyang Technological University	Singapore	SCImago #137	2


Geographic distribution of citing authors

Country	Citing papers
United States	16
China	14
United Kingdom	7
France	4
Hong Kong	3
Singapore	3
Australia	2
Spain	2
Switzerland	2
Italy	1
Mexico	1
Saudi Arabia	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	Very deep convolutional networks for large-scale image recognition	10	8 CFR 204.5(i)(3) – Outstanding Researcher
Contribution 2	Multiple View Geometry in Computer Vision	5	8 CFR 204.5(i)(3) – Outstanding Researcher
Contribution 3	The PASCAL Visual Object Classes (VOC) Challenge	10	8 CFR 204.5(i)(3) – Outstanding Researcher