

Citation Evidence Report

EB-1B Petition — Outstanding Professor or Researcher

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

Lujing Xie

Peking University

[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

99	99	4	3
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.

86.7% independent of 75 classified citing papers

Citation type	Count
Independent	65
Self-citation	0
Co-author	5
Same-institution	5

24 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 Bacon, a balanced feature-level contrastive learning framework that boosts imbalanced semi-supervised learning, establishing a novel approach to handling class imbalance in low-label regimes.

The researcher's core contribution is the development of Bacon, a method for boosting imbalanced semi-supervised learning via balanced feature-level contrastive learning, as detailed in their 2024 paper. This work stands as a distinct contribution without direct follow-up publications by the same author in the provided record.

This line of work appears to address the challenge of class imbalance within semi-supervised learning environments. By introducing balanced feature-level contrastive learning, the researcher suggests a mechanism to mitigate bias toward majority classes, offering a targeted solution for scenarios where labeled data is scarce and unevenly distributed.

The significance of this contribution is evidenced by its rapid uptake, with 22 citations recorded for the 2024 publication. Notably, 92.0% of the citing papers originate from independent researchers, indicating that the broader academic community recognizes the utility and novelty of this approach beyond the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 8 · 1 flagged influential by Semantic Scholar

CORE PAPER

[Bacon: Boosting imbalanced semi-supervised learning via balanced feature-level contrastive learning](#)

2024 · Proceedings of the AAAI Conference on Artificial Intelligence 38 (11), 11970 ..., 2024 · 22 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	A Square Peg in a Square Hole: Meta-Expert for Long-Tailed Semi-Supervised Learning	Southeast University	China	Influential
2	Lcgc: Learning from consistency gradient conflicting for class-imbalanced semi-supervised debiasing	Beijing Jiaotong University, Chongqing University of Posts and Telecommunications, Newcastle University	China, United Kingdom	—
3	Keep It on a Leash: Controllable Pseudo-label Generation Towards Realistic Long-Tailed Semi-Supervised Learning	City University of Hong Kong, Saint Francis University, Southeast University	China	—
4	A multi-view consistency framework with semi-supervised domain adaptation	Ningbo University, Zhejiang Cowain Automation Technology Co., Ltd.	China	—
5	Exploiting Minority Pseudo-Labels for Semi-Supervised Fine-Grained Road Scene Understanding	Ningbo University, Zhejiang Cowain Automation Technology Co., Ltd	China	—
6	Simple but Effective: Sub-Volume Contrastive Learning for Class-Imbalanced Semi-Supervised 3D Medical Image Segmentation	Fudan University, Guangdong University of Technology	China	—
7	ULFine: Unbiased Lightweight Fine-tuning for Foundation-Model-Assisted Long-Tailed Semi-Supervised Learning	Nanjing University of Aeronautics and Astronautics	China	—
8	Semantic Bridging Domains: Pseudo-Source as Test-Time Connector	Kuaishou Technology, Southeast University	China	—

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 rigorous benchmark for evaluating expert-level, multi-disciplinary video understanding, creating a foundational standard for assessing advanced AI capabilities in complex visual domains.

The researcher’s primary contribution centers on the 2025 publication “Mmvu: Measuring expert-level multi-discipline video understanding.” This work appears to define a critical framework for assessing how artificial intelligence systems comprehend video content across diverse, specialized fields, moving beyond general recognition to expert-level analysis.

This line of work addresses the apparent gap in standardized metrics for complex, multi-disciplinary video tasks. By focusing on “expert-level” understanding, the research suggests a shift toward more nuanced evaluation criteria that capture the depth and breadth of knowledge required in professional domains, rather than relying on superficial visual features.

The significance of this contribution is evidenced by its rapid uptake, with 98 citations recorded. Notably, 92.0% of these citations originate from independent researchers, indicating that the work has resonated broadly across the scientific community and is being utilized by external groups to advance their own studies in video understanding.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 50 · 6 flagged influential by Semantic Scholar

CORE PAPER

[Mmvu: Measuring expert-level multi-discipline video understanding](#)

2025 · Proceedings of the Computer Vision and Pattern Recognition Conference, 8475-8489, 2025 · 98 citations (GS)

Field-normalised: 103 Semantic Scholar citations place it in the top 1% of Engineering papers from 2025 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Kimi-VL Technical Report	—	—	—
2	Insight-V++: Towards Advanced Long-Chain Visual Reasoning with Multimodal Large Language Models	Nanyang Technological University, Tencent Hunyuan, Tsinghua University	China, Singapore	—
3	VLM4D: Towards Spatiotemporal Awareness in Vision Language Models	Microsoft, UCLA, University of California, Santa Cruz	United States	—
4	Vrbench: A benchmark for multi-step reasoning in long narrative videos	Nanjing University, Shanghai Artificial Intelligence Laboratory, Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences	China	—
5	Video-r1: Reinforcing video reasoning in mllms	The Chinese University of Hong Kong, The Chinese University of Hong Kong, Shenzhen, Tsinghua University	China, Hong Kong	—
6	Video-mmlu: A massive multi-discipline lecture understanding benchmark	Lambda, Inc., University of Illinois Urbana-Champaign, University of Washington	China, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
7	Wait, we don't need to "wait"! removing thinking tokens improves reasoning efficiency	University College London, University of Washington	United Kingdom, United States	—
8	Video-Holmes: Can MLLM Think Like Holmes for Complex Video Reasoning?	City University of Hong Kong, Tencent PCG	China, Hong Kong	—
9	Videorf: Incentivizing video reasoning capability in mllms via reinforced fine-tuning	Beijing Institute of Technology, Shenzhen University	China	—
10	Thinking with videos: Multimodal tool-augmented reinforcement learning for long video reasoning	ByteDance, Tsinghua Shenzhen International Graduate School, Tsinghua University, University of Chinese Academy of Sciences	China	—
11	Grpo-care: Consistency-aware reinforcement learning for multimodal reasoning	Tencent PCG, The Chinese University of Hong Kong, The University of Hong Kong	China, Hong Kong	—
12	Tinyllava-video-r1: Towards smaller lmms for video reasoning	Beihang University	China	—
13	Seed1.8 model card: Towards generalized real-world agency	ByteDance	China	—
14	Reinforcing video reasoning with focused thinking	Hefei University of Technology, Lanzhou University, National University of Singapore	China, Singapore	—
15	Video-rts: Rethinking reinforcement learning and test-time scaling for efficient and enhanced video reasoning	UNC Chapel Hill, UNC Chapel Hill, Nanyang Technological University	United States	Influential
16	Humanomniv2: From understanding to omni-modal reasoning with context	Alibaba Group	China	—
17	Judge anything: Mllm as a judge across any modality	Huazhong University of Science and Technology, University of Illinois Chicago	China, United States	—
18	Shotbench: Expert-level cinematic understanding in vision-language models	Nanyang Technological University, Shanghai Artificial Intelligence Laboratory, The Chinese University of Hong Kong	China, Singapore	—
19	Llava-critic-r1: Your critic model is secretly a strong policy model	National University of Singapore, The Ohio State University, University of Maryland, College Park	Singapore, United States	Influential
20	Onethinker: All-in-one reasoning model for image and video	CUHK, Meituan	China, Hong Kong	—
21	Avatar: Reinforcement learning to see, hear, and reason over video	Arizona State University	United States	—
22	Eoc-bench: Can mllms identify, recall, and forecast objects in an egocentric world?	Alibaba Group, Tongji University, Zhejiang University	China	—

No.	Citing paper	Citing institution(s)	Country	S2
23	When thinking drifts: Evidential grounding for robust video reasoning	The University of Texas at Austin, UC Berkeley, University of Texas at Austin	United States	—
24	Deepvideo-r1: Video reinforcement fine-tuning via difficulty-aware regressive grpo	Korea Advanced Institute of Science and Technology, Korea University	South Korea	—
25	ReWatch-R1: Boosting Complex Video Reasoning in Large Vision-Language Models through Agentic Data Synthesis	Alibaba Group	China	—
26	Videocap-r1: Enhancing mllms for video captioning via structured thinking	Honor Device Co., Ltd, Nanjing University	China	—
27	CrossVid: A Comprehensive Benchmark for Evaluating Cross-Video Reasoning in Multimodal Large Language Models	Xiaohongshu Inc.	China	—
28	SiLVR: A Simple Language-based Video Reasoning Framework	UNC Chapel Hill	United States	—
29	VKnowU: Evaluating Visual Knowledge Understanding in Multimodal LLMs	City University of Hong Kong, Nanjing University, Nanjing University, University of Science and Technology of China	China	—
30	MA-Bench: Towards Fine-grained Micro-Action Understanding	Hefei University of Technology, United Arab Emirates University, University College London	China, United Arab Emirates, United Kingdom	—

Showing the 30 most-cited of 50 independent citing papers.

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
Zhejiang University	China	SCImago #6 · THE 39 · QS 49	8
Peking University	China	SCImago #11 · THE 13 · QS 14	6
Chinese Academy of Sciences	China	SCImago #2	6
Nanjing University	China	SCImago #178 · THE =62 · QS =103	5
Yale University	United States	SCImago #76 · THE 10 · QS 21	4
National University of Singapore	Singapore	SCImago #59 · THE 17 · QS 8	4
The Chinese University of Hong Kong	Hong Kong	SCImago #163 · THE =41 · QS =32	4
City University of Hong Kong	Hong Kong	SCImago #342 · THE 73 · QS =63	4
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	4

Institution	Country	World ranking	Citing papers
University of Science and Technology of China	China	SCImago #77 · THE 51 · QS =132	4
Shanghai AI Laboratory	China	—	3
ByteDance	China	—	3
Southeast University	China	THE 251–300 · QS =392	3
Nanyang Technological University	Singapore	SCImago #137	3
New York University	United States	SCImago #116 · THE =31 · QS 55	3

Geographic distribution of citing authors

Country	Citing papers
China	58
United States	25
Singapore	7
Hong Kong	6
United Arab Emirates	4
United Kingdom	4
Canada	2
South Korea	1
Japan	1
Australia	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	Bacon: Boosting imbalanced semi-supervised learning via balanced feature-level contrastive learning	8	8 CFR 204.5(i)(3) – Outstanding Researcher
Contribution 2	Mmvu: Measuring expert-level multi-discipline video understanding	50	8 CFR 204.5(i)(3) – Outstanding Researcher