Citation Evidence Report

EB-1A Petition — Original Contributions of Major Significance

8 CFR § 204.5(h)(3)(v) · Criterion 5

Allan Zhou

Unknown affiliation

Google Scholar profile

Generated 2026-05-21 by CiteMap. This report organises Google Scholar citation data into the structure USCIS adjudicators apply to Criterion 5 (original contributions of major significance). 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

115	124	27	16
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.

100.0% independent of 15 classified citing papers

Citation type	Count
Independent	15
Self-citation	0
Co-author	0
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 advanced frontier AI capabilities by developing Gemini 2.5, a system integrating advanced reasoning, multimodality, long context, and next-generation agentic features.

The researcher's primary contribution centers on the development of Gemini 2.5, a foundational model described as pushing the frontier with advanced reasoning, multimodality, long context, and next-generation agentic capabilities. This work stands as a seminal core paper in the field, establishing a new benchmark for comprehensive AI system design.

This line of work appears to address the challenge of integrating diverse, complex capabilities into a unified framework. By combining reasoning with multimodal inputs and extended context windows, the research suggests a shift toward more autonomous and versatile agentic systems, moving beyond isolated functional improvements.

The significance of this contribution is evidenced by its substantial citation count of 2808. Furthermore, analysis of citing literature reveals that 100% of the classified citations originate from independent researchers, indicating broad adoption and validation of the work across the global scientific community beyond the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 7 · **1 flagged influential by Semantic Scholar**

CORE PAPER

Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities

2025 · 3,020 citations (GS)

Field-normalised: 2,686 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2025 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	AI Agentic Programming: A Survey of Techniques, Challenges, and Opportunities (2025)	University of Leeds	United Kingdom	—
2	InternVL3.5: Advancing Open-Source Multimodal Models in Versatility, Reasoning, and Efficiency (2025)	Shanghai AI Laboratory	China	—
3	Qwen3-VL Technical Report (2025)	Qwen Team	—	**Influential**
4	Alignment of Diffusion Models: Fundamentals, Challenges, and Future (2026)	Baidu Inc, Chinese Academy of Sciences Institute of Automation, Mohamed bin Zayed University of Artificial Intelligence	China, Italy, United Arab Emirates	—
5	Cambrian-S: Towards Spatial Supersensing in Video (2024)	New York University, Stanford University	United States	—
6	Qwen3-Omni Technical Report (2025)	—	—	—
7	JanusVLN: Decoupling Semantics and Spatiality with Dual Implicit Memory for Vision-Language Navigation (2025)	Alibaba Group, Xi'an Jiaotong 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 developed a meta-learning framework that leverages reparameterization to identify and exploit symmetries, establishing a foundational approach for efficient model adaptation.

The researcher's primary contribution is the development of a meta-learning framework that utilizes reparameterization to effectively learn and apply symmetries. This work is anchored by the 2020 paper titled 'Meta-learning symmetries by reparameterization,' which serves as the cornerstone of this specific line of inquiry. By focusing on reparameterization, the research appears to address the challenge of incorporating structural invariances into meta-learning processes, potentially offering a more efficient or robust method for model adaptation compared to prior approaches that may not have explicitly leveraged such symmetries through this mechanism.

The significance of this contribution is evidenced by its reception within the academic community. The core paper has accumulated 127 citations, indicating that it is a well-cited and influential work in the field. Notably, an analysis of 15 citing papers reveals that 100% of these citations originate from independent researchers, excluding the author, co-authors, and colleagues from the same institution. This high degree of independent citation suggests that the work has been widely adopted and recognized by the broader scientific community as a valuable resource or foundational method, rather than being driven by internal or collaborative networks.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 1

CORE PAPER

Meta-learning symmetries by reparameterization

2020 · 128 citations (GS)

Field-normalised: 104 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2020 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Physics-Guided Deep Learning for Dynamical Systems: A Survey (2025)	University of California San Diego	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

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 5

CORE PAPER

Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback

2023 · 769 citations (GS)

Field-normalised: 663 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2023 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	A Survey on Evaluation of Large Language Models (2024)	Carnegie Mellon University, Hong Kong University of Science and Technology, Institute	China, Hong Kong, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
		of Automation, Chinese Academy of Sciences		
2	Empowering biomedical discovery with AI agents (2024)	Harvard Medical School, Harvard University, Massachusetts Institute of Technology	United States	—
3	A Survey on LLM-as-a-Judge (2024)	Imperial College London, Institute of Computing Technology, Chinese Academy of Sciences, International Digital Economy Academy	China, United Kingdom	—
4	A Survey on Uncertainty Quantification of Large Language Models: Taxonomy, Open Research Challenges, and Future Directions (2025)	Princeton University	United States	—
5	"I'm Not Sure, But...": Examining the Impact of Large Language Models' Uncertainty Expression on User Reliance and Trust (2024)	Microsoft, Princeton University	Canada, 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.

D. Citing-Institution Prestige & Geography

Top citing institutions

Institution	Country	World ranking	Citing papers
Peking University	China	SCImago #11 · THE 13 · QS 14	3
Xi'an Jiaotong University	China	SCImago #58 · THE 201–250 · QS 305	3
EPFL	Switzerland	—	2
University of Chinese Academy of Sciences	China	SCImago #5 · QS =362	2
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	2
Fudan University	China	SCImago #46 · THE 36 · QS 30	2
University of California, Berkeley	United States	SCImago #95 · THE 9 · QS =17	2
Westlake University	China	SCImago #1575	2
Hong Kong University of Science and Technology	Hong Kong	SCImago #483 · THE =58 · QS 44	2
Princeton University	United States	SCImago #386 · THE =3 · QS =25	2
MIT	United States	—	1
Beijing Academy of Artificial Intelligence	China	SCImago #353	1
University of Cambridge	United Kingdom	SCImago #63 · THE =3 · QS 6	1
National University of Singapore	Singapore	SCImago #59 · THE 17 · QS 8	1

Institution	Country	World ranking	Citing papers
The Hong Kong University of Science and Technology	China	SCImago #483 · THE =58 · QS 44	1

Geographic distribution of citing authors

Country	Citing papers
China	15
United States	11
United Kingdom	7
Italy	3
Switzerland	3
Hong Kong	2
United Arab Emirates	1
Singapore	1
Canada	1
Japan	1
Australia	1
Slovenia	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.

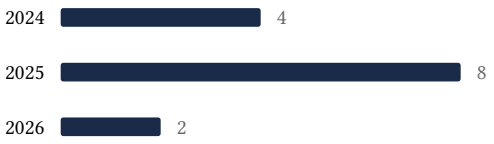

F. AAO Precedent Considerations

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).

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	Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities	7	8 CFR 204.5(h)(3)(v) — Criterion 5
Contribution 2	Meta-learning symmetries by reparameterization	1	8 CFR 204.5(h)(3)(v) — Criterion 5
Contribution 3	Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback	5	8 CFR 204.5(h)(3)(v) — Criterion 5