

Citation Evidence Report

EB-1A Petition — Original Contributions of Major Significance

8 CFR § 204.5(h)(3)(v) · Criterion 5


Haoran Li

Machine Learning Department, Carnegie Mellon University

[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



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.

88.8% independent of 152 classified citing papers

Citation type	Count
Independent	135
Self-citation	5
Co-author	12
Same-institution	0

82 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 a reliable multimodal RAG framework to enhance factuality in medical vision-language models, establishing a foundational approach for accurate clinical AI reasoning.

CLAIM: The researcher's primary contribution is the development of a reliable multimodal Retrieval-Augmented Generation (RAG) system designed to improve factuality in medical vision-language models, anchored by the 2024 paper "RULE." This work serves as the foundation for subsequent advancements in the field.

ORIGINALITY: This line of work appears to address the critical challenge of ensuring factual accuracy in medical AI systems that process both visual and textual data. The progression from the core "RULE" paper to the 2025 follow-up "MMed-RAG" suggests an evolution from establishing reliability to creating a versatile system. Additionally, the 2025 paper "HoPE" indicates parallel efforts to optimize long-context processing within these vision-language architectures, broadening the technical scope of the researcher's contributions.

SIGNIFICANCE: The impact of this research is evidenced by substantial citation counts, with the core paper accumulating 149 citations and the follow-up "MMed-RAG" reaching 197 citations. Notably, 96.1% of the 152 classified citations originate from independent researchers, demonstrating that the broader scientific community has adopted and built upon these methods outside the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 144 · 3 flagged influential by Semantic Scholar

CORE PAPER

RULE: Reliable Multimodal RAG for Factuality in Medical Vision Language Models

2024 · Conference on Empirical Methods in Natural Language Processing (EMNLP), 2024 · 149 citations (GS)

Field-normalised: 92 Semantic Scholar citations place it in the top 1% of Medicine papers from 2024 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Retrieval augmented generation and understanding in vision: A survey and new outlook	HKUST(GZ), Sichuan University, HKUST(GZ), Tsinghua University	Bulgaria, China, Italy	—
2	Mova: Adapting mixture of vision experts to multimodal context	SenseTime Research, Shanghai AI Laboratory	China, Hong Kong	—
3	Conflictbank: A benchmark for evaluating the influence of knowledge conflicts in llm	Brown University, DAMO Academy, Alibaba Group, Shanghai AI Laboratory	China, United States	—
4	Scaling inference-time search with vision value model for improved visual comprehension	Microsoft, University of Maryland, College Park	United States	—
5	Integrating Chain-of-Thought and Retrieval Augmented Generation Enhances Rare Disease Diagnosis From Clinical Notes	Children's Hospital of Philadelphia	United States	—
6	Retrieval-augmented perception: High-resolution image perception meets visual rag	Nanyang Technological University, Sun Yat-sen University, University of Sydney	Australia, China, Singapore	—

No.	Citing paper	Citing institution(s)	Country	S2
7	Retrieval-augmented generation in biomedicine: A survey of technologies, datasets, and clinical applications	Henan University, Hunan City University, LvZhiDao Information Technology Co., Ltd.	China, Singapore, Switzerland	—
8	Mira: A novel framework for fusing modalities in medical rag	Aalto University, MBZUAI	Finland, United Arab Emirates	—
9	MoMA: a mixture-of-multimodal-agents architecture for enhancing clinical prediction modelling	Northwestern University, University of Wisconsin-Madison	United States	—
10	A multimodal multi-agent framework for radiology report generation	University of North Texas	United States	—
11	: A Benchmark for Evaluating the Influence of Knowledge Conflicts in LLMs	Brown University, DAMO Academy, Alibaba Group, Shanghai AI Laboratory	China, United States	—
12	Benchmarking poisoning attacks against retrieval-augmented generation	Nankai University, University of Louisville, University of North Texas	China, United States	—
13	AlzheimerRAG: Multimodal retrieval-augmented generation for clinical use cases	Toronto Metropolitan University	Canada	—
14	Mrm: Black-box membership inference attacks against multimodal rag systems	Beijing University of Posts and Telecommunications, Tsinghua University	China	—
15	Empowering multimodal llms with external tools: A comprehensive survey	Lenovo Research, Nanyang Technological University, Xi'an Jiaotong University	China, Singapore	Influential
16	Poisoned-mrag: Knowledge poisoning attacks to multimodal retrieval augmented generation	Duke University, Huazhong University of Science and Technology, Lehigh University	China, United States	—
17	A multi-agent system for complex reasoning in radiology visual question answering	University of North Texas	United States	—
18	mrag: Elucidating the design space of multimodal retrieval-augmented generation	Texas A&M University, University of California, Berkeley, University of Texas at Austin	United States	—
19	Med-grim: enhanced zero-shot medical vqa using prompt-embedded multimodal graph rag	Shiv Nadar University	India	—
20	Who is in the Spotlight: The Hidden Bias Undermining Multimodal Retrieval-Augmented Generation	Beijing University of Posts and Telecommunications, Institute of Computing Technology, Institute of Computing Technology, Chinese Academy of Sciences	China, United States	—
21	HeteroRAG: A Heterogeneous Retrieval-Augmented Generation Framework for Medical Vision Language Tasks	Fudan University, Shanghai Jiao Tong University	China	—

No.	Citing paper	Citing institution(s)	Country	S2
22	Concept-Enhanced Multimodal RAG: Towards Interpretable and Accurate Radiology Report Generation	Sapienza University of Rome, Umeå University, UniCamillus-Saint Camillus International University of Health Sciences	Italy, Sweden	—
23	MHier-RAG: Multi-Modal RAG for Visual-Rich Document Question-Answering via Hierarchical and Multi-Granularity Reasoning	Nanjing University	China	—
24	VaccineRAG: Boosting Multimodal Large Language Models' Immunity to Harmful RAG Samples	Beihang University, QIYUAN LAB	China	—
25	AlzheimerRAG: Multimodal Retrieval Augmented Generation for Clinical Use Cases using PubMed articles	Toronto Metropolitan University	Canada	—
26	Visual Alignment of Medical Vision-Language Models for Grounded Radiology Report Generation	NEC, University of California, Riverside	United States	—
27	TemMed-Bench: Evaluating Temporal Medical Image Reasoning in Vision-Language Models	Amazon, UCLA, University of California, Los Angeles	United States	Influential
28	Safeguarding Multimodal Knowledge Copyright in the RAG-as-a-Service Environment	ShanghaiTech University, Sun Yat-sen University	China	—
29	Ask in any modality: A comprehensive survey on multimodal retrieval-augmented generation	Qatar Computing Research Institute, Sharif University of Technology, University of Tehran	Iran, Qatar	—
30	Agentpoison: Red-teaming llm agents via poisoning memory or knowledge bases	University of Chicago, University of Illinois, Urbana-Champaign, University of Wisconsin	United States	—

Showing the 30 most-cited of 61 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

MMed-RAG: Versatile Multimodal RAG System for Medical Vision Language Models

2025 · International Conference on Learning Representations (ICLR), 2025 · 197 citations (GS)

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

No.	Citing paper	Citing institution(s)	Country	S2
1	PRMBench: A Fine-grained and Challenging Benchmark for Process-Level Reward Models	—	—	—
2	Retrieval augmented generation and understanding in vision: A survey and new outlook	HKUST(GZ), HKUST(GZ), Sichuan University	Bulgaria, China, Italy	—

No.	Citing paper	Citing institution(s)	Country	S2
		University, HKUST(GZ), Tsinghua University		
3	Mova: Adapting mixture of vision experts to multimodal context	SenseTime, SenseTime Research, Shanghai AI Laboratory	China, Hong Kong	—
4	Words or vision: Do vision-language models have blind faith in text?	National University of Singapore	Singapore	—
5	Conflictbank: A benchmark for evaluating the influence of knowledge conflicts in llm	Brown University, DAMO Academy, Alibaba Group, Shanghai AI Laboratory	China, United States	—
6	Hm-rag: Hierarchical multi-agent multimodal retrieval augmented generation	Shanghai Artificial Intelligence Laboratory, The Hong Kong University of Science and Technology, The Hong Kong University of Science and Technology (Guangzhou)	China	—
7	Scaling inference-time search with vision value model for improved visual comprehension	Microsoft, University of Maryland, College Park	United States	—
8	Visa: Retrieval augmented generation with visual source attribution	CSIRO, CSIRO, University of Queensland, University of Queensland	Australia, Canada	—
9	Decalign: Hierarchical cross-modal alignment for decoupled multimodal representation learning	Texas A&M University, University of Southern California	United States	—
10	FinSage: A Multi-aspect RAG System for Financial Filings Question Answering	CG Matrix Technology Limited, Duke University, McGill University	Canada, China, United States	—
11	Patho-AgenticRAG: towards multimodal agentic retrieval-augmented generation for pathology VLMs via reinforcement learning	Independent Researcher, Shengjing Hospital of China Medical University, University of Toronto	Canada, China, New Zealand	—
12	Towards interpretable radiology report generation via concept bottlenecks using a multi-agentic RAG	CISPA Helmholtz Center for Information Security, German Research Center for Artificial Intelligence, German Research Center for Artificial Intelligence (DFKI)	Germany	—
13	A survey on agentic multimodal large language models	Chinese University of Hong Kong, City University of Hong Kong, Communication University of China	China, Singapore	—
14	End-to-end agentic RAG system training for traceable diagnostic reasoning	Shanghai Jiao Tong University, Xinhua Hospital Affiliated to Shanghai Jiao Tong University School of Medicine	China	—
15	Agentic medical knowledge graphs enhance medical question answering: Bridging the gap between llms and evolving medical knowledge	University of Toronto, Worcester Polytechnic Institute	Canada, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
16	Integrating Chain-of-Thought and Retrieval Augmented Generation Enhances Rare Disease Diagnosis From Clinical Notes	Children's Hospital of Philadelphia	United States	—
17	Retrieval-augmented perception: High-resolution image perception meets visual rag	Nanyang Technological University, Sun Yat-sen University, University of Sydney	Australia, China, Singapore	—
18	Multi-modal foundation models for computational pathology: a survey	Baylor University, Harvard Medical School, University of Arkansas	United States	—
19	Medeyes: Learning dynamic visual focus for medical progressive diagnosis	Hunan University	China	—
20	Retrieval-augmented generation in biomedicine: A survey of technologies, datasets, and clinical applications	Henan University, Hunan City University, LvZhiDao Information Technology Co., Ltd.	China, Singapore, Switzerland	—
21	Mira: A novel framework for fusing modalities in medical rag	Aalto University, MBZUAI	Finland, United Arab Emirates	—
22	MoMA: a mixture-of-multimodal-agents architecture for enhancing clinical prediction modeling	Northwestern University, University of Wisconsin-Madison	United States	—
23	Ai meets brain: Memory systems from cognitive neuroscience to autonomous agents	Fudan University, Harbin Institute of Technology, National University of Singapore	China, Singapore	—
24	A multimodal multi-agent framework for radiology report generation	University of North Texas	United States	—
25	Silent leaks: Implicit knowledge extraction attack on rag systems through benign queries	National University of Singapore, National University of Singapore, Peking University, Tsinghua University	China, Singapore	—
26	: A Benchmark for Evaluating the Influence of Knowledge Conflicts in LLMs	Brown University, DAMO Academy, Alibaba Group, Shanghai AI Laboratory	China, United States	—
27	Benchmarking poisoning attacks against retrieval-augmented generation	Nankai University, University of Louisville, University of North Texas	China, United States	—
28	AlzheimerRAG: Multimodal retrieval-augmented generation for clinical use cases	Toronto Metropolitan University	Canada	—
29	Mrm: Black-box membership inference attacks against multimodal rag systems	Beijing University of Posts and Telecommunications, Tsinghua University	China	—
30	Beyond Text: Unveiling Privacy Vulnerabilities in Multi-modal Retrieval-Augmented Generation	Amazon.com, Jilin University, Michigan State University	China, United States	—

Showing the 30 most-cited of 76 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

HoPE: Hybrid of Position Embedding for Long Context Vision-Language Models

2025 · Advances in Neural Information Processing Systems (NeurIPS), 2025 · 12 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Revisiting Multimodal Positional Encoding in Vision-Language Models	Alibaba Group, Institute of Computing Technology, Chinese Academy of Sciences	China	—
2	C^2ROPE: Causal Continuous Rotary Positional Encoding for 3D Large Multimodal-Models Reasoning	Shanghai Jiaotong University, University of Macau, University of Science and Technology of China	China	—
3	Remember Me: Bridging the Long-Range Gap in LVLMs with Three-Step Inference-Only Decay Resilience Strategies	Beijing Normal Kong Baptist University, Shanghai Jiao Tong University, University of Wollongong	Australia, China	—
4	Beyond Sequential Distance: Inter-Modal Distance Invariant Position Encoding	Chinese Academy of Sciences, Tencent	China	—
5	T2SGrid: Temporal-to-Spatial Gridification for Video Temporal Grounding	Bytedance Inc., Guangdong Laboratory of Artificial Intelligence and Digital Economy, South China University of Technology	China	—
6	No Generation without Representation: Efficient Causal Protein Language Models Enable Zero-Shot Fitness Estimation	Boston University	United States	—
7	Lumos-1: On Autoregressive Video Generation with Discrete Diffusion from a Unified Model Perspective	Alibaba Group, DAMO Academy, Alibaba Group, Tsinghua University	China, 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).

Contribution 2

Claim – Contribution 2

The researcher advanced event causality identification by developing a heuristic semantic dependency inquiry network, establishing a novel methodological framework for analyzing causal relationships in textual data.

The researcher’s contribution centers on the 2024 publication ‘Advancing Event Causality Identification via Heuristic Semantic Dependency Inquiry Network.’ This work represents a focused effort to improve the identification of causal links between events, utilizing a specific network-based approach to semantic dependencies. The title suggests a methodological innovation aimed at resolving ambiguities in how events are causally connected within complex information structures.

This line of work appears to address the challenge of accurately mapping causal relationships in text, a problem where traditional methods may lack sufficient semantic granularity. By introducing a heuristic semantic dependency inquiry network, the researcher proposed a structured way to infer causality, distinguishing this approach from simpler co-occurrence or syntactic methods. The absence of follow-up papers in the provided data indicates this stands as a singular, foundational contribution in this specific methodological niche.

The significance of this work is evidenced by its citation record, with 16 citations recorded for the core paper. Notably, within the broader context of the researcher’s portfolio, 96.1% of citing papers originate from independent researchers, suggesting that this methodological framework has been adopted and validated by the wider academic community rather than just the researcher’s immediate circle. This high degree of independent uptake underscores the utility and relevance of the proposed network approach in the field.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 2

CORE PAPER

[Advancing Event Causality Identification via Heuristic Semantic Dependency Inquiry Network](#)

2024 · Conference on Empirical Methods in Natural Language Processing (EMNLP), 2024 · 16 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	A Survey of Event Causality Identification: Taxonomy, Challenges, Assessment, and Prospects	National University of Defense Technology	China	—
2	A Survey of Event Causality Identification: Taxonomy, Challenges, Assessment, and Prospects	National University of Defense Technology	China	—

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
Shanghai Jiao Tong University	China	SCImago #10 · THE 40 · QS =47	11
Shanghai AI Laboratory	China	—	11
Nanyang Technological University	Singapore	SCImago #137	10
National University of Singapore	Singapore	SCImago #59 · THE 17 · QS 8	9
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	9
The Chinese University of Hong Kong	Hong Kong	SCImago #163 · THE =41 · QS =32	9
Zhejiang University	China	SCImago #6 · THE 39 · QS 49	8
UNC-Chapel Hill	United States	—	8
Alibaba Group	China	SCImago #226	6
Fudan University	China	SCImago #46 · THE 36 · QS 30	6
The Hong Kong University of Science and Technology	China	SCImago #483 · THE =58 · QS 44	6
Sun Yat-sen University	China	SCImago #40 · THE 201–250 · QS =276	5
University of Toronto	Canada	SCImago #39 · THE 21 · QS 29	5
University of North Texas	United States	SCImago #2445 · QS 901-950	5
City University of Hong Kong	Hong Kong	SCImago #342 · THE 73 · QS =63	5

Geographic distribution of citing authors

Country	Citing papers
China	84
United States	61
Singapore	19
Canada	14
Australia	11
United Kingdom	8
Hong Kong	6
United Arab Emirates	5
Germany	5
Japan	3
Italy	2
India	2

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	RULE: Reliable Multimodal RAG for Factuality in Medical Vision Language Models	144	8 CFR 204.5(h)(3)(v) – Criterion 5
Contribution 2	Advancing Event Causality Identification via Heuristic Semantic Dependency Inquiry Network	2	8 CFR 204.5(h)(3)(v) – Criterion 5