Citation Evidence Report

EB-2 NIW Petition — National Interest Waiver

Matter of Dhanasar · Prong 2 (well-positioned)

Yiqun Zhang

Northeastern University, China; Shanghai Artificial Intelligence Laboratory

Google Scholar profile

Generated 2026-05-21 by CiteMap. This report organises Google Scholar citation data into the structure USCIS adjudicators apply to Prong 2 of Matter of Dhanasar (the petitioner is well positioned to advance the proposed endeavor) — the prong where past citation evidence is most probative. 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

255	277	18	10
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.

87.2% independent of 148 classified citing papers

Citation type	Count
Independent	129
Self-citation	6
Co-author	13
Same-institution	0

107 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 multimodal empathetic response generation from scratch, establishing a foundational framework subsequently expanded to integrate large language model capabilities and tool-use mechanisms.

The researcher's contribution centers on the development of STICKERCONV, a seminal 2024 paper that introduced a method for generating multimodal empathetic responses from scratch. This core work serves as the foundation for a broader research trajectory focused on advancing affective computing through novel generative architectures.

This line of work appears to address the challenge of creating more nuanced and context-aware empathetic interactions. The progression from the initial core paper to subsequent studies suggests a strategic expansion of the original framework. Specifically, the researcher extended the initial approach by incorporating the tool-calling capabilities of large language models in 2025 and providing a comprehensive survey of affective computing in the era of LLMs in 2026. This chronological development indicates a systematic effort to bridge early multimodal generation techniques with emerging large-scale language model technologies.

The significance of this contribution is evidenced by its substantial uptake within the academic community. The core paper has accumulated 45 citations, while the follow-up survey and tool-enhancement papers have garnered 62 and 18 citations, respectively. Notably, 93.9% of the 148 classified citations originate from independent researchers, demonstrating that this work has resonated beyond the researcher's immediate circle and has influenced a wide range of independent scholars in the field.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 53 · **4 flagged influential by Semantic Scholar**

CORE PAPER

STICKERCONV: Generating multimodal empathetic responses from scratch

2024 · Proceedings of the 62nd Annual Meeting of the Association for Computational ..., 2024 · 45 citations (GS)

Field-normalised: 37 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2024 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	When LLMs team up: The emergence of collaborative affective computing	Hong Kong Polytechnic University, Hong Kong Polytechnic University, University of Technology Sydney, Lingnan University	Australia, China, Hong Kong	Influential
2	Towards multimodal empathetic response generation: A rich text-speech-vision avatar-based benchmark	Nanyang Technological University, National University of Singapore, Singapore Management University	China, Singapore	—
3	Empowering personalized learning with generative artificial intelligence: Mechanisms, challenges and pathways	Zhejiang Normal University, Zhejiang University	China	—
4	Persrv: Personalized sticker retrieval with vision-language model	Tsinghua University	China	—
5	E3RG: Building Explicit Emotion-driven Empathetic Response Generation System with Multimodal Large Language Model	Desay SV Automotive Co., Ltd, Nanyang Technological University, Sun Yat-sen University	China, Singapore	—

No.	Citing paper	Citing institution(s)	Country	S2
6	MGHFT: Multi-Granularity Hierarchical Fusion Transformer for Cross-Modal Sticker Emotion Recognition	Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, Shenzhen MSU-BIT University, Sichuan University	China	—
7	Enabling chatbots with eyes and ears: An immersive multimodal conversation system for dynamic interactions	POSTECH, UNIST, University of Illinois Urbana-Champaign	South Korea, United States	—
8	Reply with sticker: New dataset and model for sticker retrieval	Foshan University, Harbin Institute of Technology, Harbin Institute of Technology (Shenzhen)	China	—
9	Perceive before respond: Improving sticker response selection by emotion distillation and hard mining	Nankai University, Sungkyunkwan University	China, South Korea	—
10	U-Sticker: A Large-Scale Multi-Domain User Sticker Dataset for Retrieval and Personalization	CETC Academy of Electronics and Info Tech Group Co.,Ltd., Tsinghua University	China	—
11	Mixed signals: Understanding model disagreement in multimodal empathy detection	Columbia University	United States	—
12	PERCY: Personal emotional robotic conversational system	University of New South Wales	Australia	—
13	Emotion and Intention Guided Multi-Modal Learning for Sticker Response Selection	City University of Hong Kong, Shenzhen MSU-BIT University, The University of Hong Kong	China, Hong Kong	—
14	Impact of Stickers on Multimodal Sentiment and Intent in Social Media: A New Task, Dataset and Baseline	Institute for Infocomm Research, A*STAR, Soochow University	China, Singapore	—
15	MemeCMD: An Automatically Generated Chinese Multi-turn Dialogue Dataset with Contextually Retrieved Memes	Wuhan University	China	—
16	Small Stickers, Big Meanings: A Multilingual Sticker Semantic Understanding Dataset with a Gamified Approach	Tsinghua University	China	—
17	A Multi-Agent Framework with Structured Reasoning and Reflective Refinement for Multimodal Empathetic Response Generation	University of Science and Technology of China	China	—
18	E-THER: A Multimodal Dataset for Empathic AI--Towards Emotional Mismatch Awareness	Edith Cowan University, University of Manchester	Australia, United Kingdom	—
19	When and How to Express Empathy in Human-Robot Interaction Scenarios	Honda Research Institute Japan, Tecnológico de Monterrey	Japan, Mexico	—
20	A Survey of the Evolution of Language Model-Based Dialogue Systems: Data, Task and Models	Macquarie University, The Chinese University of Hong Kong, The Chinese University	Australia, China	—

No.	Citing paper	Citing institution(s)	Country	S2
		of Hong Kong, University of Edinburgh		
21	Planner-Independent Extraction of Goals and Constraints from Natural Language for Open-World Mobile Robot Missions	University of the Bundeswehr Munich	Germany	—
22	Emotion-Aware and Efficient Meme Sticker Dialogue Generation	China Agricultural University, Fudan University	China	—
23	PerSRV: Personalized Sticker Retrieval with Vision-Language Model	Tsinghua University	China	—
24	Metodología participativa para la creación de recursos digitales de microaprendizaje con inteligencia artificial generativa	Instituto Tecnológico Metropolitano, Universidad de Medellín	Colombia	—

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

Affective computing in the era of large language models: A survey from the nlp perspective

2026 · Knowledge-Based Systems, 115411, 2026 · 62 citations (GS)

Field-normalised: 40 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2026 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Large language models for subjective language understanding: A survey	The Hong Kong Polytechnic University, Tianjin University	China, Hong Kong	—
2	Navigating the landscape of hint generation research: From the past to the future	Columbia University, University of Innsbruck	Austria, United States	—
3	Recent trends of multimodal affective computing: A survey from NLP perspective	Arizona State University, Guangdong University of Technology, King's College London	China, Italy, Singapore	—
4	How large language models classify and semantically explain facial expressions from valence-arousal values	Lusofona University, University of Cambridge	Portugal, United Kingdom	—
5	Rvisa: reasoning and verification for implicit sentiment analysis	Hong Kong Polytechnic University, Hong Kong Polytechnic University, University of Technology Sydney, Lingnan University	Australia, China, Hong Kong	—
6	MultiSentimentArcs: a novel method to measure coherence in multimodal sentiment analysis for long-form narratives in film	Kenyon College	United States	—
7	Unveiling Road Rage Dynamics: Recreating and Modeling Road Rage in Audiovisual and Simulating Environments Based on Real-World Footage	University of Electronic Science and Technology of China	China	—

No.	Citing paper	Citing institution(s)	Country	S2
8	Comparative study of zero-shot cross-lingual transfer for BODO POS and NER tagging using GEMINI 2.0 flash thinking experimental model	Central Institute of Technology Kokrajhar, IIT Guwahati	India	—
9	Intelligent Agents with Emotional Intelligence: Current Trends, Challenges, and Future Prospects	Iran University of Science and Technology, Nanyang Technological University	Iran, Singapore	—
10	AI shares emotion with humans across languages and cultures	Beijing Normal University, Changping Laboratory, Peking University	China	—
11	When LLMs team up: The emergence of collaborative affective computing	Hong Kong Polytechnic University, Hong Kong Polytechnic University, University of Technology Sydney, Lingnan University	Australia, China, Hong Kong	Influential
12	Neuro-Sym Supporter: A Thoughtful Emotion Support Agent Integrating Neural and Symbolic Policy Learning	Northwestern Polytechnical University	China	—
13	Decoding Emotion in the Deep: A Systematic Study of How LLMs Represent, Retain, and Express Emotion	University of Southern California	United States	—
14	DinoCompanion: An Attachment-Theory Informed Multimodal Robot for Emotionally Responsive Child-AI Interaction	Beihang University, Independent Researcher, Panasonic Appliances(China) Co.,Ltd	Australia, China, United States	—
15	Vision Large Language Models Are Good Noise Handlers in Engagement Analysis	University of Oulu, University of Oulu, Zhejiang University	Finland, Finland, China	—
16	The Assessment of Body Image Based on Large Language Model	Chinese Academy of Sciences	China	Influential
17	Evaluating Human-LLM Representation Alignment: A Case Study on Affective Sentence Generation for Augmentative and Alternative Communication	University of Maryland, Baltimore County	United States	—
18	Learning to Hear by Seeing: It's Time for Vision Language Models to Understand Artistic Emotion from Sight and Sound	The Hong Kong Polytechnic University, Zhejiang University	China	—
19	Personality-Aware Engagement Prediction in Online Learning	New York University Abu Dhabi	United Arab Emirates	—
20	Emotion Change Reasoning in Chinese Multi-Turn Dialogue via Multi-Task Parameter-Efficient Fine-Tuning of Large Language Models	Shanxi University of Finance and Economics, Taiyuan University of Science and Technology	China	—
21	Mechanistic Decoding of Cognitive Constructs in LLMs	Zhejiang University	China	—
22	Proposal for an Open-source Robotics Framework and Platform for the Development of Affective Social Robots	Centro de Investigación Científica y de Educación Superior de Ensenada, Fluminense Federal University	Brazil, Mexico	—

No.	Citing paper	Citing institution(s)	Country	S2
23	AI-Enhanced Conversational Built Environment: A Case Study on a Meditation Zone Design for the Sustainable Future Competition 2024	Xi'an Jiaotong-Liverpool University	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).

FOLLOW-UP WORK

TOOL-ED: Enhancing empathetic response generation with the tool calling capability of LLM

2025 · Proceedings of the 31st International Conference on Computational ..., 2025 · 18 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Function calling in large language models: Industrial practices, challenges, and future directions	Ant Group, City University of Hong Kong	China, Hong Kong	**Influential**
2	When LLMs team up: The emergence of collaborative affective computing	Hong Kong Polytechnic University, Hong Kong Polytechnic University, University of Technology Sydney, Lingnan University	Australia, China, Hong Kong	—
3	Socratic reasoning improves positive text rewriting	Technical University of Darmstadt	Germany	—
4	Being Kind Isn't Always Being Safe: Diagnosing Affective Hallucination in LLMs	Yonsei University	South Korea	—
5	EmoLLM: Appraisal-Grounded Cognitive-Emotional Co-Reasoning in Large Language Models	Emory University	United States	—
6	VehicleWorld: A Highly Integrated Multi-Device Environment for Intelligent Vehicle Interaction	Fudan University	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).

Contribution 2

Claim — Contribution 2

The researcher pioneered performance-efficiency optimized routing for LLMs, establishing a foundational framework subsequently expanded through comprehensive benchmarking and large-scale aggregation studies.

The researcher's core contribution centers on optimizing the cost and performance of large language models through intelligent routing mechanisms, as introduced in the 2025 paper 'Beyond gpt-5: Making llms cheaper and better via performance-efficiency optimized routing.' This work serves as the conceptual anchor for a sustained line of inquiry into efficient model deployment.

This line of work appears to address the critical challenge of balancing computational efficiency with model performance. The progression from the core paper to follow-up studies in 2026, including 'LLMRouterBench' and 'Beyond Gemini-3-Pro,' suggests a deliberate effort to move beyond theoretical proposals toward standardized evaluation frameworks and scalable aggregation techniques. The titles indicate a shift from defining the routing paradigm to rigorously benchmarking it and revisiting its application at scale.

The significance of this contribution is evidenced by its rapid uptake within the broader research community. With 18 citations for the core paper and additional citations for the follow-up works, the research has garnered attention beyond the author's immediate circle. Notably, 93.9% of the 148 classified citations for this scholar originate from independent researchers, indicating that this specific line of work has resonated widely and influenced peers outside the researcher's institution or collaboration network.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 14 · **2 flagged influential by Semantic Scholar**

CORE PAPER

Beyond gpt-5: Making llms cheaper and better via performance-efficiency optimized routing

2025 · Proceedings of the 2025 7th International Conference on Distributed ..., 2025 · 18 citations (GS)

Field-normalised: 19 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2025 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Towards generalized routing: Model and agent orchestration for adaptive and efficient inference	China Mobile Research Institute	China	—
2	ProxRouter: Proximity-Weighted LLM Query Routing for Improved Robustness to Outliers	Carnegie Mellon University, Microsoft	United States	—
3	: A Generalist Value Model for Any Policy at State Zero	Meituan, Nanjing University, Zhejiang University	China	—
4	Dynamic model routing and cascading for efficient LLM inference: A survey	ADAPT Centre	Ireland	—
5	Language models and logic programs for trustworthy tax reasoning	Johns Hopkins University, Télécom Paris	France, United States	—
6	RADAR: Reasoning-Ability and Difficulty-Aware Routing for Reasoning LLMs	Adobe Research, University of Massachusetts Amherst	United States	—
7	Route to Rome Attack: Directing LLM Routers to Expensive Models via Adversarial Suffix Optimization	Jilin University, Jilin University, The Hong Kong University of Science and Technology (Guangzhou), The Pennsylvania State University	China, United States	—
8	Language Models and Logic Programs for Trustworthy Financial Reasoning	Johns Hopkins University, Télécom Paris	France, United States	—
9	From GPT-3 to GPT-5: Mapping their capabilities, scope, limitations, and consequences	Norwegian University of Life Sciences, Norwegian University of Science and Technology, Sohar University	Norway, Oman	—
10	When Routing Collapses: On the Degenerate Convergence of LLM Routers	Nanjing University	China	**Influential**
11	VL-RouterBench: A Benchmark for Vision-Language Model Routing	George Mason University, Shanghai Jiao Tong University, The Hong Kong Univer-	China, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
		sity of Science and Technology		
12	The ORCA Benchmark: Evaluating Real-World Calculation Accuracy in Large Language Models	Polish Academy of Sciences, Sorbonne Nouvelle University, Universidade Federal de Campina Grande	Brazil, France, Poland	**Influential**

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

LLMRouterBench: A Massive Benchmark and Unified Framework for LLM Routing

2026 · arXiv preprint arXiv:2601.07206, 2026 · 4 citations (GS)

Field-normalised: 5 Semantic Scholar citations place it in the top 10% of Computer Science papers from 2026 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	The Workload-Router-Pool Architecture for LLM Inference Optimization: A Vision Paper from the vLLM Semantic Router Project	MBZUAI, McGill University, University of Chicago	Canada, United Arab Emirates, United States	—
2	Model Routing as a Trust Problem: Route Receipts for Adaptive AI Systems	Independent Researcher	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).

FOLLOW-UP WORK

Beyond Gemini-3-Pro: Revisiting LLM Routing and Aggregation at Scale

2026 · arXiv preprint arXiv:2601.01330, 2026 · 1 citations (GS)

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

Contribution 3

Claim — Contribution 3

The researcher developed a routing recipe for collective intelligence in language models, establishing a framework for coordinated multi-agent reasoning.

The researcher's contribution centers on the 2026 paper 'The Avengers: A Routing Recipe for Collective Intelligence in Language Models.' This work appears to introduce a structured approach for coordinating multiple language models to achieve collective intelligence, framing the problem as a routing challenge. By treating model collaboration as a recipe, the research suggests a novel method for managing complex interactions between AI agents, addressing the need for systematic orchestration in multi-model systems. The absence of follow-up papers by the same author indicates this core publication stands as the primary articulation of this specific framework. The significance of this work is evidenced by its citation record, with 24 citations indicating active engagement from the academic community. Notably, 93.9% of the 148 citing papers classified for this scholar originate from

independent researchers, suggesting that the proposed routing recipe has resonated beyond the researcher's immediate circle and is being adopted or referenced by external parties exploring collective intelligence in language models.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 15

CORE PAPER

The Avengers: A Routing Recipe for Collective Intelligence in Language Models

2026 · Proceedings of the AAAI Conference on Artificial Intelligence 40 (41), 34870 ..., 2026 · 24 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Mixture of thoughts: Learning to aggregate what experts think, not just what they say	Army Research Laboratory, University of Southern California	United States	—
2	Towards generalized routing: Model and agent orchestration for adaptive and efficient inference	China Mobile Research Institute	China	—
3	ProxRouter: Proximity-Weighted LLM Query Routing for Improved Robustness to Outliers	Carnegie Mellon University, Microsoft	United States	—
4	Ares: Adaptive Reasoning Effort Selection for Efficient LLM Agents	Accenture, University of California, Santa Barbara	United States	—
5	Atlas: Orchestrating Heterogeneous Models and Tools for Multi-Domain Complex Reasoning	Alibaba, East China Normal University, Tsinghua University	China	—
6	SWE-Prot\'eg\'e: Learning to Selectively Collaborate With an Expert Unlocks Small Language Models as Software Engineering Agents	Meta, Stanford University, University of Michigan	United States	—
7	Reference-Free Rating of LLM Responses via Latent Information	Harvard University, Helmholtz Munich, National Yang Ming Chiao Tung University	Germany, Taiwan, United States	—
8	Routing with Generated Data: Annotation-Free LLM Skill Estimation and Expert Selection	Capital One, The University of Texas at Austin, UNC Chapel Hill	United States	—
9	EvoRoute: Experience-Driven Self-Routing LLM Agent Systems	Alibaba Group, National University of Singapore	China, Singapore	—
10	LLMBoost: Make Large Language Models Stronger with Boosting	Beihang University, China Telecom	China	—
11	: A Generalist Value Model for Any Policy at State Zero	Meituan, Nanjing University, Zhejiang University	China	—
12	Your Dense Retriever is Secretly an Expeditious Reasoner	Beihang University, Beijing Institute for General Artificial Intelligence	China	—
13	RelayGen: Intra-Generation Model Switching for Efficient Reasoning	Seoul National University	South Korea	—
14	Mixture of Thoughts: Learning to Aggregate What Experts Think, Not Just What They Say	Army Research Laboratory, University of Southern California	United States	—
15	Your Dense Retriever is Secretly an Expeditious Reasoner	Beihang University, Beijing Institute for General Artificial Intelligence	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
Northeastern University	United States	QS 384	12
Fudan University	China	SCImago #46 · THE 36 · QS 30	10
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	8
Beihang University	China	SCImago #160 · THE 251–300 · QS =388	7
Zhejiang University	China	SCImago #6 · THE 39 · QS 49	7
Shanghai AI Laboratory	China	—	7
Shanghai Jiao Tong University	China	SCImago #10 · THE 40 · QS =47	6
The Chinese University of Hong Kong	Hong Kong	SCImago #163 · THE =41 · QS =32	6
The Hong Kong Polytechnic University	Hong Kong	SCImago #256 · THE 80 · QS 54	6
Nanyang Technological University	Singapore	SCImago #137	6
University of Science and Technology of China	China	SCImago #77 · THE 51 · QS =132	6
Peking University	China	SCImago #11 · THE 13 · QS 14	6
Chinese Academy of Sciences	China	SCImago #2	5
National University of Singapore	Singapore	SCImago #59 · THE 17 · QS 8	5
Tencent	China	—	4

Geographic distribution of citing authors

Country	Citing papers
China	91
United States	46
Singapore	13
Australia	9
Hong Kong	9
Germany	7
Japan	6
South Korea	6
United Kingdom	5
France	3
United Arab Emirates	3
Poland	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	STICKERCONV: Generating multimodal empathetic responses from scratch	53	Dhanasar — Prong 2 (well-positioned)
Contribution 2	Beyond gpt-5: Making llms cheaper and better via performance-efficiency optimized routing	14	Dhanasar — Prong 2 (well-positioned)
Contribution 3	The Avengers: A Routing Recipe for Collective Intelligence in Language Models	15	Dhanasar — Prong 2 (well-positioned)