

Citation Evidence Report

EB-1B Petition — Outstanding Professor or Researcher

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


Fanheng Kong

Northeastern University; Kuaishou Technology

[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.

97.1% independent of 207 classified citing papers

Citation type	Count
Independent	201
Self-citation	1
Co-author	5
Same-institution	0

252 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 methods for generating multimodal empathetic responses from scratch, establishing a foundational framework subsequently extended to fine-grained temporal understanding in dynamic video contexts.

The researcher's contribution centers on the development of novel approaches for generating multimodal empathetic responses from scratch, as demonstrated in the core paper "STICKERCONV" (2024). This work serves as the foundation for a broader research trajectory that includes subsequent investigations into comprehensive fine-grained temporal understanding, such as the 2025 paper "TUNA," which evaluates dense dynamic videos. This progression suggests a strategic expansion from static or turn-based empathetic generation to complex, time-sensitive multimodal analysis.

The originality of this line of work appears to lie in addressing the challenge of synthesizing empathetic responses without relying on pre-existing templates, thereby enabling more authentic and context-aware interactions. The follow-up work on temporal understanding indicates an effort to refine these capabilities for dynamic, real-world video data, suggesting a methodological evolution from basic response generation to sophisticated, time-aware evaluation frameworks.

The significance of this research is evidenced by the substantial uptake of the core paper, which has accumulated 45 citations. Notably, 99.0% of the citing papers originate from independent researchers, indicating that the work has resonated widely across the broader academic community rather than within a single institutional circle. This high degree of independent citation underscores the foundational impact and broad relevance of the researcher's contributions to multimodal AI.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 29 · 2 flagged influential by Semantic Scholar

CORE PAPER

STICKERCONV: Generating multimodal empathetic responses from scratch

2024 · ACL2024, 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

[TUNA: Comprehensive Fine-grained Temporal Understanding Evaluation on Dense Dynamic Videos](#)

2025 · ACL2025, 2025 · 6 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Hud: Hierarchical uncertainty-aware disambiguation network for composed video retrieval	Harbin Institute of Technology (Shenzhen), Shandong University	China	—
2	Visual jigsaw post-training improves mllms	Linköping University, Nanyang Technological University, SenseTime Research	Singapore, Sweden	—
3	Learning Transferable Temporal Primitives for Video Reasoning via Synthetic Videos	Alibaba Group, Tsinghua University, Zhejiang University	China	—
4	OmniJigsaw: Enhancing Omni-Modal Reasoning via Modality-Orchestrated Reordering	Xiaomi Inc., Zhejiang University	China	Influential
5	Can Vision-Language Models Solve the Shell Game?	National University of Singapore	Singapore	—

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 critically evaluated the efficacy of Mamba architectures for time series forecasting, establishing a foundational benchmark that has garnered significant independent scholarly attention.

The researcher's contribution centers on the critical assessment of Mamba models within the domain of time series forecasting, as detailed in the 2025 paper titled 'Is mamba effective for time series forecasting?' This work serves as the primary anchor for this line of inquiry, standing alone without direct follow-up publications by the same author in the provided dataset.

This line of work appears to address a timely gap in understanding the practical utility of state-space models like Mamba for sequential data prediction. By posing a direct evaluative question, the research likely provided necessary empirical clarity or theoretical grounding regarding the suitability of these architectures for forecasting tasks, distinguishing itself through its focused and diagnostic approach.

The significance of this contribution is evidenced by its substantial citation count of 376, indicating rapid uptake within the field. Notably, 99.0% of the citing papers originate from independent researchers, suggesting that the work has resonated broadly across the academic community and influenced external research directions rather than merely circulating within a single group.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 147 · 14 flagged influential by Semantic Scholar

CORE PAPER

[Is mamba effective for time series forecasting?](#)

2025 · Neurocomputing 619, 129178, 2025 · 376 citations (GS)

Field-normalised: 263 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	A comprehensive survey of deep learning for time series forecasting: architectural diversity and open challenges	Seoul National University	South Korea	—
2	A survey on time-series pre-trained models	South China University of Technology, The Hong Kong University of Science and Technology	China	Background
3	Unlocking the power of lstm for long term time series forecasting	Alibaba Group, Duke Kunshan University, Princeton University	China, United Kingdom, United States	—
4	Time-o1: Time-series forecasting needs transformed label alignment	Peking University, Renmin University of China, Xiaohongshu Inc.	China	—
5	Bi-mamba+: Bidirectional mamba for time series forecasting	Beijing University of Posts and Telecommunications, China Telecom Corporation, China Telecom Research Institute	China	—
6	Decision mamba: A multi-grained state space model with self-evolution regularization for offline rl	Great Bay University, Harbin Institute of Technology (Shenzhen)	China	—
7	xlstm-mixer: Multivariate time series forecasting by mixing via scalar memories	Eindhoven University of Technology, TU Darmstadt	Germany, Netherlands	Influential
8	Decomposed spatio-temporal Mamba for long-term traffic prediction	Beijing University of Technology	China	—
9	Attractor memory for long-term time series forecasting: A chaos perspective	Griffith University, Squirrel Ai Learning, The Hong Kong University of Science and Technology (Guangzhou)	Australia, China	Background

No.	Citing paper	Citing institution(s)	Country	S2
10	Block-biased mamba for long-range sequence processing	University of Pittsburgh	United States	—
11	Avs-mamba: Exploring temporal and multi-modal mamba for audio-visual segmentation	Dalian University of Technology	China	—
12	TSCMamba: Mamba meets multi-view learning for time series classification	University of Kentucky	United States	—
13	SST: Multi-Scale Hybrid Mamba-Transformer Experts for Time Series Forecasting	Emory University, Illinois Institute of Technology, Northwestern University	United States	—
14	CMMamba: channel mixing Mamba for time series forecasting	Xinjiang University	China	—
15	Wavelet mixture of experts for time series forecasting	Ningbo University, Shanghai University of Engineering Science	China	—
16	Madiff: Motion-aware mamba diffusion models for hand trajectory prediction on egocentric videos	National University of Defense Technology, Shanghai Jiao Tong University	China	Influential
17	Time-ssm: Simplifying and unifying state space models for time series forecasting	Hong Kong University of Science and Technology (Guangzhou), The Hong Kong University of Science and Technology (Guangzhou)	China	Background
18	Mamba meets financial markets: A graph-mamba approach for stock price prediction	Simon Fraser University, The University of British Columbia, University of Pittsburgh	Canada, United States	—
19	Ehrmamba: Towards generalizable and scalable foundation models for electronic health records	Vector Institute, Vector Institute; University of Toronto	Canada	—
20	SSD-TS: Exploring the potential of linear state space models for diffusion models in time series imputation	East China Normal University	China	—
21	MI-Mamba: A hybrid motor imagery electroencephalograph classification model with Mamba's global scanning	Tianjin University	China	—
22	Beyond sensor data: Foundation models of behavioral data from wearables improve health predictions	Apple, Apple Inc.	United States	—
23	Mamba4cast: Efficient zero-shot time series forecasting with state space models	ELLIS Institute Tübingen & University of Freiburg, University of Freiburg	Germany	—
24	Fmamba: Mamba based on fast-attention for multivariate time-series forecasting	University of Science and Technology of China, USTC	China	—
25	Exploring neural granger causality with xL-STMs: Unveiling temporal dependencies in complex data	Carnegie Mellon University, Eindhoven University of Technology, TU Darmstadt	Germany, Netherlands, United States	—
26	Mambular: A sequential model for tabular deep learning	BASF, Clausthal University of Technology, LMU Munich	Canada, Germany	—

No.	Citing paper	Citing institution(s)	Country	S2
27	A novel state space model with dynamic graph neural network for EEG event detection	Fudan University, Shanghai Jiao Tong University	China	—
28	CMDMamba: dual-layer Mamba architecture with dual convolutional feed-forward networks for efficient financial time series forecasting	Guangxi Police College, Guangxi Vocational and Technical Institute of Industry	China	—
29	Integrated spatio-temporal modeling with hybrid graph convolutions and the graph fourier neural operator for traffic prediction	Ferdowsi University of Mashhad	Iran	—
30	Technologies on effectiveness and efficiency: A survey of state spaces models	Carnegie Mellon University, Tsinghua University	China, United States	Influential

Showing the 30 most-cited of 147 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).

Contribution 3

Claim – Contribution 3

The researcher developed a modality curation framework for building universal embeddings to advance multimodal information retrieval systems.

The researcher's contribution centers on the 2025 paper 'Modality curation: Building universal embeddings for advanced multimodal information retrieval.' This work appears to establish a foundational approach for creating universal embeddings that facilitate advanced retrieval across multiple data modalities.

This line of work addresses the challenge of integrating diverse data types into a unified retrieval framework. By focusing on modality curation, the research suggests a novel method for constructing embeddings that are broadly applicable, potentially overcoming limitations in existing multimodal systems that struggle with cross-modal alignment.

The significance of this contribution is evidenced by its rapid uptake, with 22 citations recorded. Notably, 99.0% of the 207 citing papers classified for this scholar originate from independent researchers, indicating that the broader academic community recognizes the utility and originality of this approach in advancing multimodal information retrieval.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 19 · 2 flagged influential by Semantic Scholar

CORE PAPER

[Modality curation: Building universal embeddings for advanced multimodal information retrieval](#)

2025 · arXiv preprint arXiv:2505.19650, 2025 · 22 citations (GS)

Field-normalised: 20 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	Intent: Invariance and discrimination-aware noise mitigation for robust composed image retrieval	Shandong University	China	—

No.	Citing paper	Citing institution(s)	Country	S2
2	Habit: Chrono-synergia robust progressive learning framework for composed image retrieval	Shandong University	China	—
3	Retrack: Evidence-driven dual-stream directional anchor calibration network for composed video retrieval	Shandong Jianzhu University, Shandong University	China	—
4	Hud: Hierarchical uncertainty-aware disambiguation network for composed video retrieval	Harbin Institute of Technology (Shenzhen), Shandong University	China	—
5	Embed-rl: Reinforcement learning for reasoning-driven multimodal embeddings	Kuaishou Technology, Tsinghua Shenzhen International Graduate School, Tsinghua University, Tsinghua University	China	—
6	Tembed: Unlocking task scaling in universal multimodal embeddings	OPPO Research Institute, University of Macau	China	—
7	Rzenembed: Towards comprehensive multimodal retrieval	360 AI Research	—	—
8	From generator to embedder: Harnessing innate abilities of multimodal llms via building zero-shot discriminative embedding model	Korea University	South Korea	—
9	Metaembed: Scaling multimodal retrieval at test-time with flexible late interaction	Meta, Rice University	United States	—
10	On the role of pretrained language models in general-purpose text embeddings: A survey	Harbin Institute of Technology	China	—
11	Beyond Chain-of-Thought: Rewrite as a Universal Interface for Generative Multimodal Embeddings	Hefei Comprehensive National Science Center, Tencent Inc., Tsinghua University	China	—
12	Towards Universal Video Retrieval: Generalizing Video Embedding via Synthesized Multimodal Pyramid Curriculum	Alibaba Group, HKUST(GZ)	China	Influential
13	Combating Visual Neglect and Semantic Drift in Large Multimodal Models for Enhanced Cross-Modal Retrieval	Baidu Inc.	China	—
14	Hierarchical Long Video Understanding with Audiovisual Entity Cohesion and Agentic Search	Microsoft Research Asia, University of Science and Technology of China	China	—
15	ReMatch: Boosting Representation through Matching for Multimodal Retrieval	Huazhong University of Science and Technology, University of Glasgow, Xiaohongshu Inc.	China, United Kingdom	—
16	PREGEN: Uncovering Latent Thoughts in Composed Video Retrieval	Ben-Gurion University of the Negev, Bosch Center for AI	Israel	Influential
17	Generative Giants, Retrieval Weaklings: Why do Multimodal Large Language Models Fail at Multimodal Retrieval?	Peking University, University of Electronic Science and Technology of China	China	—

No.	Citing paper	Citing institution(s)	Country	S2
18	Reasoning Guided Embeddings: Leveraging MLLM Reasoning for Improved Multimodal Retrieval	Beijing Institute of Technology, Nanjing University, SenseTime Research	China	—
19	Let Multimodal Embedders Learn When to Augment Query via Adaptive Query Augmentation	—	—	—

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
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	15
Fudan University	China	SCImago #46 · THE 36 · QS 30	8
Zhejiang University	China	SCImago #6 · THE 39 · QS 49	7
Nanyang Technological University	Singapore	SCImago #137	6
University of Science and Technology of China	China	SCImago #77 · THE 51 · QS =132	5
Shandong University	China	SCImago #79 · THE 251–300 · QS =339	5
The University of Hong Kong	Hong Kong	SCImago #195 · THE 33 · QS 11	5
Peking University	China	SCImago #11 · THE 13 · QS 14	4
Shanghai Jiao Tong University	China	SCImago #10 · THE 40 · QS =47	4
Yale University	United States	SCImago #76 · THE 10 · QS 21	4
Wuhan University	China	SCImago #80 · THE =122 · QS 186	4
The Hong Kong University of Science and Technology (Guangzhou)	China	SCImago #483 · THE =58 · QS 44	4
University of Hong Kong	China	SCImago #195 · THE 33 · QS 11	4
Northeastern University	United States	QS 384	4
Alibaba Group	China	SCImago #226	4

Geographic distribution of citing authors

Country	Citing papers
China	131
United States	42
Canada	11
Germany	10
South Korea	10
Australia	9

Country	Citing papers
Singapore	9
United Kingdom	9
Hong Kong	7
Japan	4
Taiwan	3
United Arab Emirates	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	29	8 CFR 204.5(i)(3) – Outstanding Researcher
Contribution 2	Is mamba effective for time series forecasting?	147	8 CFR 204.5(i)(3) – Outstanding Researcher
Contribution 3	Modality curation: Building universal embeddings for advanced multimodal information retrieval	19	8 CFR 204.5(i)(3) – Outstanding Researcher