Citation Evidence Report

EB-2 NIW Petition — National Interest Waiver

Matter of Dhanasar · Prong 2 (well-positioned)

Ming Wang

Ph.D. student of Singapore Management University & Northeastern University

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

549	555	18	7
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.

97.6% independent of 246 classified citing papers

Citation type	Count
Independent	240
Self-citation	3
Co-author	3
Same-institution	0

303 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 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 2025 publication titled 'Is mamba effective for time series forecasting?', which serves as the core work in this line of inquiry. This paper appears to address the emerging question of whether Mamba-based state space models offer practical advantages for temporal data prediction, a topic of growing interest in machine learning research.

By posing this specific evaluative question, the work likely fills a gap in understanding the applicability of novel sequence modeling architectures to time series domains. The title suggests a critical assessment rather than a mere application, indicating an original contribution to the methodological discourse on model selection for forecasting tasks.

The significance of this work is evidenced by its citation record, with 379 citations indicating substantial uptake by the broader academic community. Notably, 97.6% of these citations originate from independent researchers, demonstrating that the findings have resonated beyond the author's immediate circle and have influenced independent scholarly efforts in the field.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 148 · **14 flagged influential by Semantic Scholar**

CORE PAPER

Is mamba effective for time series forecasting?

2025 · Neurocomputing 619, 129178, 2025 · 379 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

No.	Citing paper	Citing institution(s)	Country	S2
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
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	—

No.	Citing paper	Citing institution(s)	Country	S2
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 xLSTMs: 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	—
27	A novel state space model with dynamic graphic 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 148 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 2

Claim — Contribution 2

The researcher developed LangGPT, a structured reusable prompt design framework for LLMs that rethinks prompt engineering through the lens of programming language principles.

The researcher's core contribution is the development of LangGPT, a framework that rethinks structured reusable prompt design for large language models by applying concepts from programming languages. This work is anchored in the 2024 paper titled 'LangGPT: Rethinking structured reusable prompt design framework for LLMs from the programming language.'

This line of work appears to address the need for more systematic and reusable approaches to prompt engineering. By drawing parallels with programming languages, the researcher suggests a novel method for structuring prompts, moving beyond ad-hoc techniques to a more formalized framework. The absence of follow-up papers indicates this seminal work stands alone as the primary vehicle for this specific contribution.

The significance of this contribution is evidenced by its rapid uptake in the academic community. With 37 citations, the paper has garnered attention, and notably, 97.6% of the 246 citing papers classified for this scholar are from independent researchers. This high degree of citation independence suggests that the LangGPT framework has resonated broadly across the field, influencing work by researchers outside the author's immediate circle and institution.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 19

CORE PAPER

LangGPT: Rethinking structured reusable prompt design framework for LLMs from the programming language

2024 · arXiv preprint arXiv:2402.16929, 2024 · 37 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Flipattack: Jailbreak llms via flipping	National University of Singapore	Singapore	—
2	Mathfusion: Enhancing mathematical problem-solving of llm through instruction fusion	Renmin University of China, Shanghai Artificial Intelligence Laboratory, Shanghai Jiao Tong University	China	—
3	From prompts to templates: A systematic prompt template analysis for real-world LLMapps	Technical University of Munich	Germany	—
4	Beyond prompt content: Enhancing LLM performance via content-format integrated prompt optimization	Fudan University, Microsoft Research Asia	China	—
5	Omgm: Orchestrate multiple granularities and modalities for efficient multimodal retrieval	Microsoft Research Asia	China	—
6	Can generative AI effectively perform quality evaluation within social sciences? A case study in library and information science	Nanjing University	China	—
7	PromptPrism: A linguistically-inspired taxonomy for prompts	Amazon AWS	—	—
8	Prompt orchestration markup language	Microsoft Research	United States	—
9	Can GPT tell us why these images are synthesized? Empowering Multimodal Large Language Models for Forensics	Chinese Academy of Sciences, Institute of Information Engineering, Chinese Academy of Sciences	China	—
10	An intelligent generative method of fashion design combining attribute knowledge and Stable Diffusion Model	East China University of Science and Technology	China	—
11	Revealing the Unseen: AI Chain on LLMs for Predicting Implicit Dataflows to Generate Dataflow Graphs in Dynamically Typed Code	CSIRO, Jiangxi Normal University	Australia, China	—
12	Open Local Knowledge Graph Construction from Academic Papers Using Generative Large Language Models	Australian National University	Australia	—
13	A Chinese language learning application that integrates generative AI into augmented reality environment to improve children's literacy and reading skills	Central China Normal University, Jingzhou Vocational College of Technology	China	—
14	Does language bias GenAI academic evaluation in humanities and social sciences? A mixed-methods study based on Chinese-language HSS papers	Nanjing University	China	—
15	Processing unstructured clinical notes with LLMs: applying the CMQOE framework for hypertension	China University of Petroleum (East China), The Second Affiliated Hospital of Army Medical University (Xinqiao Hospital), The Third Affiliated Hospital of CQMU(FangDa Hospital)	China	—

No.	Citing paper	Citing institution(s)	Country	S2
16	DelvePO: Direction-Guided Self-Evolving Framework for Flexible Prompt Optimization	Nanjing University	China	—
17	OOPrompt: Reifying Intents into Structured Artifacts for Modular and Iterative Prompting	University of California, Los Angeles	United States	—
18	The PICCO Framework for Large Language Model Prompting: A Taxonomy and Reference Architecture for Prompt Structure	Mayo Clinic College of Medicine and Science	United States	—
19	Code Retrieval with Mixture of Experts Prototype Learning Based on Classification	Guangdong University of Technology, Huizhou University, Macau Polytechnic 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 3

Claim — Contribution 3

The researcher developed MM-InstructEval, a framework for zero-shot evaluation of multimodal large language models on reasoning tasks, establishing a benchmark for assessing model capabilities without task-specific training.

The researcher's contribution centers on the development of MM-InstructEval, a seminal work published in 2025 that addresses the evaluation of multimodal large language models. This core paper introduces a methodology for zero-shot assessment of these models on multimodal reasoning tasks, providing a standardized approach to measuring performance without the need for task-specific fine-tuning.

This line of work appears to address a critical gap in the field by offering a robust mechanism for evaluating the reasoning capabilities of multimodal systems in a zero-shot setting. By focusing on instruction-based evaluation, the research suggests a shift towards more generalizable and efficient assessment protocols, distinguishing itself from prior methods that may have relied on extensive training data or narrow task definitions.

The significance of this contribution is evidenced by its rapid uptake within the academic community, with the core paper accumulating 61 citations. Notably, the broader citation context for this scholar reveals that 97.6% of citing papers originate from independent researchers, indicating that this work has resonated widely beyond the researcher's immediate circle and has become a recognized reference point for independent studies in multimodal AI evaluation.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 32 · **1 flagged influential by Semantic Scholar**

CORE PAPER

MM-InstructEval: Zero-Shot Evaluation of (Multimodal) Large Language Models on Multimodal Reasoning Tasks

2025 · Information Fusion 122, 103204, 2025 · 61 citations (GS)

Field-normalised: 28 Semantic Scholar citations place it in the top 10% of Computer Science papers from 2025 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Large language models meet text-centric multimodal sentiment analysis: A survey	Chinese Academy of Sciences, Harbin Institute of Technology	China	—

No.	Citing paper	Citing institution(s)	Country	S2
2	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	—
3	Aligning vision to language: Annotation-free multimodal knowledge graph construction for enhanced llms reasoning	East China Normal University, New York University, Shanghai Artificial Intelligence Laboratory	China, United States	—
4	A survey of multimodal sarcasm detection	George Mason University, IITB-Monash Research Academy, Lancaster University	India, United Kingdom, United States	—
5	Multimodal large language models meet multimodal emotion recognition and reasoning: A survey	Central South University of Forestry and Technology, State University of New York New Paltz	China, United States	—
6	Memverse: Multimodal memory for lifelong learning agents	Shanghai AI Laboratory, Shanghai Artificial Intelligence Laboratory	China	—
7	Multimodal chain of continuous thought for latent-space reasoning in vision-language models	Harvard University	United States	—
8	The future of mllm prompting is adaptive: A comprehensive experimental evaluation of prompt engineering methods for robust multimodal performance	University College Dublin, University of California San Diego	Ireland, United States	—
9	Can Large Vision-Language Models Understand Multimodal Sarcasm?	The University of Texas at Dallas	United States	—
10	Evaluating open-source vision-language models for multimodal sarcasm detection	George Mason University, IITB-Monash Research Academy, Lancaster University	India, United Kingdom, United States	—
11	Deqa: Descriptions enhanced question-answering framework for multimodal aspect-based sentiment analysis	JD AI Research, Nankai University	China	—
12	MER-Bench: A Comprehensive Benchmark for Multimodal Meme Reappraisal	Anhui University, Hefei University of Technology, United Arab Emirates University	China, United Arab Emirates	—
13	MMIFEvol: Towards Evolutionary Multimodal Instruction Following	Fudan University, Huawei	China	—
14	Hit-RAG: Learning to Reason with Long Contexts via Preference Alignment	Huazhong University of Science and Technology, Shanghai Artificial Intelligence Laboratory, Shenzhen University of Advanced Technology	Australia, China, United States	—
15	Bitdecoding: Unlocking tensor cores for long-context llms with low-bit kv cache	Microsoft Research, University of Edinburgh	United Kingdom, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
16	Beyond single frames: Can lmms comprehend temporal and contextual narratives in image sequences?	Peking University, The Hong Kong Polytechnic University, Tsinghua University	China	—
17	CCAF: Coarse-to-fine Cross-Modal Alignment and Fusion for Multimodal Sentiment Analysis	Harbin Institute of Technology, Harbin Institute of Technology, Shenzhen	China	—
18	UAVBench and UAVIT-1M: Benchmarking and Enhancing MLLMs for Low-Altitude UAV Vision-Language Understanding	Northwestern Polytechnical University	China	—
19	Learning in order! a sequential strategy to learn invariant features for multimodal sentiment analysis	Harbin Institute of Technology, Monash University	Australia, China	—
20	More-r1: Guiding lvlm for multimodal object-entity relation extraction via stepwise reasoning with reinforcement learning	AlignBase, Peking University	China	—
21	Evaluating multimodal large language models on spoken sarcasm understanding	University of Groningen	Netherlands	—
22	SilVar: Speech-Driven Multimodal Model for Reasoning Visual Question Answering and Object Localization	FPT University, Harvard University, University of Alabama at Birmingham	Sweden, United States, Vietnam	—
23	Social Caption: Evaluating Social Understanding in Multimodal Models	Carnegie Mellon University	United States	—
24	Towards Multimodal Sentiment Analysis Via Contrastive Cross-Modal Retrieval Augmentation and Hierachical Prompts	Chinese People's Liberation Army General Hospital, Harbin Institute of Technology, Jiangnan University	China	—
25	Decoding hate in memes: multimodal and multitask approaches for low-resource Indonesian social media	Muhammadiyah University of Surakarta	Indonesia	—
26	A Diffusion Driven Multimodal Fusion Framework for Context Aware Sarcasm Detection via Sentiment Syntax Graph Modeling	C. V. Raman Global University, Veer Surendra Sai University of Technology	India	—
27	Customizing Visual Emotion Evaluation for MLLMs: An Open-vocabulary, Multifaceted, and Scalable Approach	Institute of Information Engineering, Chinese Academy of Sciences, Nankai University, Tsinghua University	China	—
28	Stepwise Schema-Guided Prompting Framework with Parameter Efficient Instruction Tuning for Multimedia Event Extraction	01.AI, Baidu Inc., Peking University	China	—
29	A Novel MLLMs-Based Two-Stage Model for Zero-Shot Multimodal Sentiment Analysis	Jiangnan University	China	—
30	Cognitive Hives: A Distributed Systems Theory for Conflict-Aware Multimodal Intelligence	Independent Researcher	United States	**Influential**

Showing the 30 most-cited of 32 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).

D. Citing-Institution Prestige & Geography

Top citing institutions

Institution	Country	World ranking	Citing papers
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	13
Peking University	China	SCImago #11 · THE 13 · QS 14	11
Fudan University	China	SCImago #46 · THE 36 · QS 30	11
Harbin Institute of Technology	China	SCImago #56 · THE =131 · QS 256	7
Shanghai Artificial Intelligence Laboratory	China	SCImago #563	6
Zhejiang University	China	SCImago #6 · THE 39 · QS 49	6
The Hong Kong University of Science and Technology (Guangzhou)	China	SCImago #483 · THE =58 · QS 44	6
Shanghai Jiao Tong University	China	SCImago #10 · THE 40 · QS =47	6
Northeastern University	United States	QS 384	5
Tongji University	China	SCImago #82 · THE =141 · QS =177	5
Harvard University	United States	SCImago #4 · THE =5 · QS 5	4
Yale University	United States	SCImago #76 · THE 10 · QS 21	4
Nanyang Technological University	Singapore	SCImago #137	4
University of Science and Technology of China	China	SCImago #77 · THE 51 · QS =132	4
University of Hong Kong	China	SCImago #195 · THE 33 · QS 11	4

Geographic distribution of citing authors

Country	Citing papers
China	148
United States	57
Canada	13
Germany	12
United Kingdom	11
Australia	10
Singapore	7
South Korea	7
Hong Kong	7
Japan	5
United Arab Emirates	4
Switzerland	4

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	Is mamba effective for time series forecasting?	148	Dhanasar — Prong 2 (well-positioned)
Contribution 2	LangGPT: Rethinking structured reusable prompt design framework for LLMs from the programming language	19	Dhanasar — Prong 2 (well-positioned)
Contribution 3	MM-InstructEval: Zero-Shot Evaluation of (Multimodal) Large Language Models on Multimodal Reasoning Tasks	32	Dhanasar — Prong 2 (well-positioned)