

Citation Evidence Report

EB-1A Petition — Original Contributions of Major Significance

8 CFR § 204.5(h)(3)(v) · Criterion 5


Zihan Wang

Northeastern University (China)

[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.

97.8% independent of 413 classified citing papers

Citation type	Count
Independent	404
Self-citation	3
Co-author	6
Same-institution	0

335 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 deep search agents using knowledge graphs and multi-turn RL, establishing a framework subsequently adapted for specialized medical AI applications.

The researcher's core contribution centers on the 2025 paper 'Deepdive,' which advances deep search agents through the integration of knowledge graphs and multi-turn reinforcement learning. This work serves as the foundational pillar for a subsequent line of inquiry focused on specialized agentic systems.

Originality in this trajectory is suggested by the evolution from general deep search to domain-specific applications. The titles of follow-up works, such as 'Deepmed' and 'Medxiaohu,' indicate a deliberate extension of the core methodology into the medical field. This progression implies a novel approach to building medical multi-modal large language models and research agents via multi-hop search and turn-controlled training, leveraging the architectural insights established in the initial study.

The significance of this research line is evidenced by its uptake within the broader academic community. With 36 citations for the core paper and additional citations for the follow-up works, the methodology has attracted attention. Notably, the vast majority of citations across the researcher's portfolio originate from independent researchers, suggesting that the proposed frameworks for deep search and medical AI agents are being adopted and built upon by peers outside the researcher's immediate institution or collaboration network.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 30 · 5 flagged influential by Semantic Scholar

CORE PAPER

[Deepdive: Advancing deep search agents with knowledge graphs and multi-turn rl](#)

2025 · arXiv preprint arXiv:2509.10446, 2025 · 36 citations (GS)

Field-normalised: 31 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	Nested browser-use learning for agentic information seeking	Alibaba Group	China	Influential
2	Beyond turn limits: Training deep search agents with dynamic context window	Alibaba Group, Chinese Academy of Sciences, City University of Hong Kong	China	—
3	Infoagent: Advancing autonomous information-seeking agents	Brown University, Microsoft, Microsoft Research Asia	China, United States	Influential
4	Openseeker: Democratizing frontier search agents by fully open-sourcing training data	Shanghai Jiao Tong University	China	Influential
5	AgentFold: Long-Horizon Web Agents with Proactive Context Management	Alibaba Group, Shanghai Jiao Tong University	China	Influential
6	Browsemaster: Towards scalable web browsing via tool-augmented programmatic agent pair	Shanghai Jiao Tong University	China	—
7	Agentic Aggregation for Parallel Scaling of Long-Horizon Agentic Tasks	Princeton University	United States	—
8	Websailor-v2: Bridging the chasm to proprietary agents via synthetic data and scalable reinforcement learning	Alibaba Group, Shanghai Jiao Tong University	China	—

No.	Citing paper	Citing institution(s)	Country	S2
9	Chaining the Evidence: Robust Reinforcement Learning for Deep Search Agents with Citation-Aware Rubric Rewards	Tsinghua University, Zhipu AI	China	Influential
10	AgentFold: Long-Horizon Web Agents with Proactive Context Folding	Alibaba Group, Shanghai Jiao Tong University	China	—
11	Towards Scalable Web Browsing via Tool-Augmented Programmatic Agent Pair	Shanghai Jiao Tong University	China	—
12	Function calling in large language models: Industrial practices, challenges, and future directions	Ant Group, City University of Hong Kong	China, Hong Kong	—
13	Webleaper: Empowering efficiency and efficacy in webagent via enabling info-rich seeking	Alibaba Group	China	—
14	To Search or Not to Search: Aligning the Decision Boundary of Deep Search Agents via Causal Intervention	City University of Hong Kong, Huawei Technologies Ltd.	China, Hong Kong	—
15	Agentrl: Scaling agentic reinforcement learning with a multi-turn, multi-task framework	Tsinghua University	China	—
16	SAGE: Steerable Agentic Data Generation for Deep Search with Execution Feedback	Google, Google DeepMind, New York University	United Kingdom, United States	—
17	Beyond pipelines: A survey of the paradigm shift toward model-native agentic ai	Beijing Jiaotong University	China	—
18	Parallelmuse: Agentic parallel thinking for deep information seeking	Alibaba Group	China	—
19	Marco DeepResearch: Unlocking Efficient Deep Research Agents via Verification-Centric Design	Alibaba International Digital Commerce	China	—
20	OffSeeker: Online Reinforcement Learning Is Not All You Need for Deep Research Agents	Fudan University, Tencent Hunyuan, The University of Hong Kong	China, Singapore	—
21	RE-TRAC: REcursive TRAjectory Compression for Deep Search Agents	Brown University, Microsoft, Microsoft Research Asia	China, United States	—
22	Unlocking the Power of Multi-Agent LLM for Reasoning: From Lazy Agents to Deliberation	Michigan State University, Microsoft, Sea AI Lab	United States	—
23	GraphScout: Empowering Large Language Models with Intrinsic Exploration Ability for Agentic Graph Reasoning	Hangzhou City University, Nanyang Technological University, Zhejiang University	China, Singapore	—
24	MMDeepResearch-Bench: A Benchmark for Multimodal Deep Research Agents	Amazon, Case Western Reserve University, CUHK	China, Hong Kong, United Kingdom	—
25	Synthesizing Agentic Data for Web Agents with Progressive Difficulty Enhancement Mechanisms	Salesforce AI Research, University of Texas at Austin, University of Wisconsin-Madison	United States	—
26	Mind DeepResearch Technical Report	Li Auto Inc	China	—

No.	Citing paper	Citing institution(s)	Country	S2
27	WideSeek: Advancing Wide Research via Multi-Agent Scaling	Chinese Academy of Sciences, University of Science and Technology of China	China	—
28	Synthesizing Agentic Data for Web Agent Training with Progressive Difficulty Enhancement	Salesforce AI Research, University of Texas at Austin, University of Wisconsin-Madison	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

[Deepmed: Building a medical deepresearch agent via multi-hop med-search data and turn-controlled agentic training & inference](#)

2026 · arXiv preprint arXiv:2601.18496, 2026 · 3 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	TheraAgent: Multi-Agent Framework with Self-Evolving Memory and Evidence-Calibrated Reasoning for PET Theranostics	Fudan University, Nanjing University, National University of Singapore	China, Singapore, Switzerland	—

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

[Medxiaohe: A comprehensive recipe for building medical mllms](#)

2026 · arXiv preprint arXiv:2602.12705, 2026 · 2 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	TheraAgent: Multi-Agent Framework with Self-Evolving Memory and Evidence-Calibrated Reasoning for PET Theranostics	Fudan University, Nanjing University, National University of Singapore	China, Singapore, Switzerland	—

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 effectiveness 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 models, a novel class of state space models, offer practical advantages for temporal data prediction compared to established methods. By

posing this direct evaluative question, the work likely provides a critical empirical assessment that helps define the utility and limitations of this architecture in the forecasting domain.

The originality of this contribution lies in its timely examination of a rapidly evolving model class. As Mamba architectures gained traction, there was a need for rigorous validation in specific applications like time series. The researcher’s work appears to fill this gap by providing a focused analysis, thereby guiding subsequent research directions and model selection strategies within the community. The absence of follow-up papers by the same researcher suggests this single study stands as a definitive, self-contained assessment of the topic.

The significance of this work is evidenced by its substantial citation count of 376, indicating rapid and widespread uptake by the academic community. Notably, 99.3% of the citing papers originate from independent researchers, demonstrating that the findings have resonated beyond the author’s immediate circle. This high degree of independent citation suggests the work has become a standard reference point for scholars investigating state space models in time series analysis, validating its impact on the field.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 145 · 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	Bi-mamba+: Bidirectional mamba for time series forecasting	Beijing University of Posts and Telecommunications, China Telecom Corporation, China Telecom Research Institute	China	—
5	Decision mamba: A multi-grained state space model with self-evolution regularization for offline rl	Great Bay University, Harbin Institute of Technology (Shenzhen)	China	—
6	xlstm-mixer: Multivariate time series forecasting by mixing via scalar memories	Eindhoven University of Technology, TU Darmstadt	Germany, Netherlands	Influential
7	Decomposed spatio-temporal Mamba for long-term traffic prediction	Beijing University of Technology	China	—
8	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
9	Block-biased mamba for long-range sequence processing	University of Pittsburgh	United States	—
10	Avs-mamba: Exploring temporal and multi-modal mamba for audio-visual segmentation	Dalian University of Technology	China	—
11	TSCMamba: Mamba meets multi-view learning for time series classification	University of Kentucky	United States	—
12	SST: Multi-Scale Hybrid Mamba-Transformer Experts for Time Series Forecasting	Emory University, Illinois Institute of Technology, Northwestern University	United States	—
13	CMMamba: channel mixing Mamba for time series forecasting	Xinjiang University	China	—
14	Wavelet mixture of experts for time series forecasting	Ningbo University, Shanghai University of Engineering Science	China	—
15	Madiff: Motion-aware mamba diffusion models for hand trajectory prediction on egocentric videos	National University of Defense Technology, Shanghai Jiao Tong University	China	Influential
16	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
17	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	—
18	Ehrmamba: Towards generalizable and scalable foundation models for electronic health records	Vector Institute, Vector Institute; University of Toronto	Canada	—
19	SSD-TS: Exploring the potential of linear state space models for diffusion models in time series imputation	East China Normal University	China	—
20	MI-Mamba: A hybrid motor imagery electroencephalograph classification model with Mamba's global scanning	Tianjin University	China	—
21	Beyond sensor data: Foundation models of behavioral data from wearables improve health predictions	Apple, Apple Inc.	United States	—
22	Mamba4cast: Efficient zero-shot time series forecasting with state space models	ELLIS Institute Tübingen & University of Freiburg, University of Freiburg	Germany	—
23	Fmamba: Mamba based on fast-attention for multivariate time-series forecasting	University of Science and Technology of China, USTC	China	—
24	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	—
25	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
26	A novel state space model with dynamic graphic neural network for EEG event detection	Fudan University, Shanghai Jiao Tong University	China	—
27	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	—
28	Integrated spatio-temporal modeling with hybrid graph convolutions and the graph fourier neural operator for traffic prediction	Ferdowsi University of Mashhad	Iran	—
29	Technologies on effectiveness and efficiency: A survey of state spaces models	Carnegie Mellon University, Tsinghua University	China, United States	Influential
30	Autoformer: Efficient hierarchical autoregressive transformer for time series prediction	Aalborg University, Cambridge University, National University of Singapore	China, Denmark, Hong Kong	—

Showing the 30 most-cited of 145 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 token and sequence-level reward shaping methods incorporating policy entropy to enhance reinforcement learning training stability and efficiency.

The researcher's core contribution is articulated in the 2025 paper 'Gtpo and grpo-s: Token and sequence-level reward shaping with policy entropy.' This work appears to introduce novel mechanisms for adjusting reward signals at both token and sequence levels, specifically integrating policy entropy to refine the training process of language models or similar sequential decision-making systems.

This line of work addresses the challenge of stabilizing reinforcement learning objectives by leveraging entropy as a regularization or shaping factor. By distinguishing between token-level and sequence-level adjustments, the research suggests a more granular approach to reward design, potentially mitigating issues such as reward hacking or training instability that are common in large-scale model optimization.

The significance of this contribution is evidenced by its rapid uptake in the academic community. With 32 citations in a short timeframe, and notably 99.3% of citing papers originating from independent researchers, the work demonstrates broad external validation. This high degree of independent citation indicates that the methodology has been adopted and built upon by the wider scientific community, underscoring its practical utility and theoretical impact.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 22 · 2 flagged influential by Semantic Scholar

CORE PAPER

[Gtpo and grpo-s: Token and sequence-level reward shaping with policy entropy](#)

2025 · arXiv preprint arXiv:2508.04349, 2025 · 32 citations (GS)

Field-normalised: 34 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	Rethinking entropy interventions in rlvr: An entropy change perspective	Independent Researcher, Tencent, Zhejiang University	China, United States	—
2	Espo: Entropy importance sampling policy optimization	Shopee Pte. Ltd.	Singapore	—
3	How to allocate, how to learn? dynamic roll-out allocation and advantage modulation for policy optimization	Fudan University, Meituan, Peking University	China	Influential
4	Densegrpo: From sparse to dense reward for flow matching model alignment	Alibaba Group, Huazhong University of Science and Technology	China	—
5	d-TreeRPO: Towards More Reliable Policy Optimization for Diffusion Language Models	Alibaba Group, Peking University, Tsinghua University	China, United States	—
6	Skip-Connected Policy Optimization for Implicit Advantage	Carnegie Mellon University, Mohamed bin Zayed University of Artificial Intelligence, Renmin University of China	China, United Arab Emirates, United States	—
7	CLIPO: Contrastive Learning in Policy Optimization Generalizes RLVR	Alibaba, Chinese Academy of Sciences	China	—
8	Demystifying Design Choices of Reinforcement Fine-tuning: A Batched Contextual Bandit Learning Perspective	Daqing Oilfield Chongqing Company, iFLYTEK, University of Science and Technology of China	China	—
9	WS-GRPO: Weakly-Supervised Group-Relative Policy Optimization for Rollout-Efficient Reasoning	Adobe Research, The University of New South Wales, UCSD	Australia, Canada, United States	—
10	FROST: Filtering Reasoning Outliers with Attention for Efficient Reasoning	Iowa State University, Northwestern University, RTX Technology Research Center	United States	—
11	Rubrics to Tokens: Bridging Response-level Rubrics and Token-level Rewards in Instruction Following Tasks	Alibaba Group, Shanghai Jiao Tong University, Zhejiang University	China	—
12	Rethinking Reinforcement fine-tuning of LLMs: A Multi-armed Bandit Learning Perspective	iFLYTEK, University of Science and Technology of China	China	—
13	Policy Split: Incentivizing Dual-Mode Exploration in LLM Reinforcement with Dual-Mode Entropy Regularization	Beihang University, Beijing Institute of Technology, ByteDance	China	—
14	The Role of Entropy in Visual Grounding: Analysis and Optimization	Fudan University, Hikvision	China	—
15	Towards Generalizable Reasoning: Group Causal Counterfactual Policy Optimization for LLM Reasoning	Institute of Software Chinese Academy of Sciences; University of the Chinese Academy of Sciences, Peking University, The Hong Kong University of Science and Technology	China	—
16	Targeted Exploration via Unified Entropy Control for Reinforcement Learning	Chinese Academy of Sciences, Nankai University; Zhong-	China	—

No.	Citing paper	Citing institution(s)	Country	S2
		guancun Academy, Shanghai Jiao Tong University		
17	Distribution-Centric Policy Optimization Dominates Exploration-Exploitation Trade-off	Beijing Institute of Technology, Nankai University, Zhejiang University	China	—
18	SHAPE: Stage-aware Hierarchical Advantage via Potential Estimation for LLM Reasoning	Huawei, Peking University, Shanghai University of Finance and Economics	China	—
19	Orchestrating Tokens and Sequences: Dynamic Hybrid Policy Optimization for RLVR	Shopee Pte. Ltd., Tsinghua University, Xiamen University	China, Singapore	—
20	Triviality Corrected Endogenous Reward	Alibaba Group, Lanzhou University, Peking University	China	Influential
21	Reverse Browser: Vector-Image-to-Code Generator	University of Oxford	United Kingdom	—
22	Robust Object Detection for Autonomous Driving via Curriculum-Guided Group Relative Policy Optimization	Shandong Normal 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).

D. Citing-Institution Prestige & Geography

Top citing institutions

Institution	Country	World ranking	Citing papers
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	37
Shanghai Jiao Tong University	China	SCImago #10 · THE 40 · QS =47	33
Alibaba Group	China	SCImago #226	30
Peking University	China	SCImago #11 · THE 13 · QS 14	23
Fudan University	China	SCImago #46 · THE 36 · QS 30	22
Zhejiang University	China	SCImago #6 · THE 39 · QS 49	21
University of Science and Technology of China	China	SCImago #77 · THE 51 · QS =132	16
Chinese Academy of Sciences	China	SCImago #2	14
The University of Hong Kong	Hong Kong	SCImago #195 · THE 33 · QS 11	11
Nanyang Technological University	Singapore	SCImago #137	10
Meituan	China	—	10
Beijing University of Posts and Telecommunications	China	SCImago #355 · QS 1001-1200	10
National University of Singapore	Singapore	SCImago #59 · THE 17 · QS 8	9
Tencent	United States	—	9
Hong Kong University of Science and Technology	Hong Kong	SCImago #483 · THE =58 · QS 44	8

Geographic distribution of citing authors

Country	Citing papers
China	286
United States	95
Hong Kong	25
Singapore	24
Canada	22
United Kingdom	17
South Korea	13
Australia	11
Germany	10
Japan	6
Switzerland	6
United Arab Emirates	5

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	Deepdive: Advancing deep search agents with knowledge graphs and multi-turn rl	30	8 CFR 204.5(h)(3)(v) – Criterion 5
Contribution 2	Is mamba effective for time series forecasting?	145	8 CFR 204.5(h)(3)(v) – Criterion 5
Contribution 3	Gtpo and grpo-s: Token and sequence-level reward shaping with policy entropy	22	8 CFR 204.5(h)(3)(v) – Criterion 5