

Citation Evidence Report

EB-2 NIW Petition — National Interest Waiver

Matter of Dhanasar · Prong 2 (well-positioned)


Yingxin

Graph Origin

[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



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.

95.8% independent of 333 classified citing papers

Citation type	Count
Independent	319
Self-citation	0
Co-author	14
Same-institution	0

164 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 a segmentation-based approach to deepfake detection and localization, subsequently expanding this framework to address multimodal evaluation and generalized multi-scenario robustness.

The researcher's core contribution rests on the 2023 paper 'Detect any deepfakes: Segment anything meets face forgery detection and localization,' which appears to integrate segmentation techniques with forgery detection. This work established a foundational method for identifying and localizing facial forgeries, achieving 55 citations.

This line of work addresses the evolving complexity of synthetic media by moving beyond basic detection. The titles of follow-up papers suggest a strategic expansion: 'Shield' (2025) introduces a benchmark for multimodal large language models in face spoofing, while 'GM-DF' (2025) focuses on generalized multi-scenario detection. This progression indicates an effort to enhance robustness and evaluation standards in the field.

The significance of this research is underscored by its uptake within the academic community. The core paper has garnered 55 citations, while the 2025 follow-ups have accumulated 69 and 38 citations respectively. Notably, 100% of the 333 classified citations for this scholar originate from independent researchers, demonstrating broad external validation and impact beyond the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 91 · 3 flagged influential by Semantic Scholar

CORE PAPER

[Detect any deepfakes: Segment anything meets face forgery detection and localization](#)

2023 · Chinese conference on biometric recognition, 180-190, 2023 · 55 citations (GS)

Field-normalised: 34 Semantic Scholar citations place it in the top 10% of Computer Science papers from 2023 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	IKA2: Internal Knowledge Adaptive Activation for Robust Recognition in Complex Scenarios	Beijing Jiaotong University, Jingtangshan University, Tianjin University of Science and Technology	China	—
2	Fantastic animals and where to find them: Segment any marine animal with dual sam	Dalian University of Technology	China	—
3	Credit card fraud detection using advanced transformer model	Johns Hopkins University, Miami University, Northeastern University	United States	—
4	Unmasking synthetic realities in generative AI: A comprehensive review of adversarially robust deepfake detection systems	Hamad Bin Khalifa University	Qatar	—
5	Application of adaptive machine learning systems in heterogeneous data environments	Elevance Health, Georgia Institute of Technology, Google LLC	United States	—
6	Opensdi: Spotting diffusion-generated images in the open world	Harbin Institute of Technology, University of Southampton, Xi'an Jiaotong University	China, United Kingdom	—
7	Evaluating modern approaches in 3d scene reconstruction: Nerf vs gaussian-based methods	Columbia University, Cornell University, Guangxi University	China, Germany, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
8	Knowledge graph embedding and few-shot relational learning methods for digital assets in usa	Canoakbit Alliance Inc, Emory University	United States	—
9	Mapping new realities: Ground truth image creation with pix2pix image-to-image translation	Saarland University of Applied Sciences, Southern Methodist University, Texas A&M University	Germany, United States	Methodology
10	Image forgery classification and localization through vision transformers	University of Hyderabad, Vardhaman College of Engineering	India	—
11	Fintech and MSEs innovation: An empirical analysis	Renmin University of China	China	Background
12	Adapi: Facilitating dnn model adaptivity for efficient private inference in edge computing	NC State University, Northeastern University, UMass Dartmouth	United States	—
13	Mindev: Multi-modal integrated diffusion framework for video reconstruction from eeg signals	Bytedance Inc., Shanghai Jiao Tong University, University of Shanghai for Science and Technology	China	—
14	Delocate: detection and localization for deepfake videos with randomly-located tampered traces	Hunan University, Nanyang Technological University, National University of Singapore	China, Singapore	Background
15	Segment anything for videos: A systematic survey	Chinese Academy of Sciences, Shanghai Jiao Tong University, The Hong Kong University of Science and Technology (Guangzhou)	China, Finland	Methodology
16	Evaluating SAM2 for Video Semantic Segmentation	Fudan University, Guizhou University, Institute for Information Research, A*STAR	China, Singapore	—
17	A novel universal image forensics localization model based on image noise and segment anything model	Shenzhen University	China	—
18	HFP-SAM: Hierarchical Frequency Prompted SAM for Efficient Marine Animal Segmentation	Affiliated Zhongshan Hospital of Dalian University, Central Hospital of Dalian University of Technology, Dalian University of Technology	China	—
19	A noise and edge extraction-based dual-branch method for shallowfake and deepfake localization	Delhi Technological University	India	—
20	Patch-Discontinuity Mining for Generalized Deepfake Detection	Chinese Academy of Military Science, Shanghai Jiao Tong University	China	—
21	Boosting non-causal semantic elimination: An unconventional harnessing of LVM for open-world deepfake interpretation	Beijing Jiaotong University, Lenovo Research	China	—

No.	Citing paper	Citing institution(s)	Country	S2
22	Frequency Domain Enhanced U-Net for Low-Frequency Information-Rich Image Segmentation in Surgical and Deep-Sea Exploration Robots	Peking University, University of Chinese Academy of Sciences, University of Electronic Science and Technology of China	China	—
23	Shallowfake and deepfake image manipulation localization using noise and RGB-based dual branch method	Delhi Technological University	India	—
24	APGNet: Adaptive Prior-Guided for Underwater Camouflaged Object Detection	Nanjing University of Aeronautics and Astronautics, University of Leicester	China, United Kingdom	—
25	Hybrid deep learning architecture for comprehensive deepfake video facial forgery detection and segmentation	University of Hyderabad	India	—
26	Ethical and Legal Concerns of Deepfake Technology in Biomedical Imaging: A Comprehensive Survey: S. Rajpar et al.	University of Science and Technology Beijing	China	—
27	LSP-ST: Ladder Shape-Biased Side-Tuning for Robust Infrared Small Target Detection	Changchun Institute of Optics, Fine Mechanics and Physics, Chinese Academy of Sciences, Sun Yat-sen University	China	—
28	Semantic segmentation of marine animal image by U-Net based on multi-cognitive visual adapter and dual-attention fusion mechanism	China University of Petroleum (East China)	China	—
29	Research on the integration of deep learning and psychology in intelligent digital education technology	Cachon University, Dankook University, Hefei Changjiang Roud Kindergarten Education Group Jingfu Branch School	China, South Korea	—
30	FTCNet: A Foreground Transformation Contrast Network for Marine Animal Segmentation	Qilu University of Technology (Shandong Academy of Sciences), Shandong University, Shenzhen Technology University	China	—

Showing the 30 most-cited of 31 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 is Influential signal, Valenzuela et al. 2015), or *Background* (a passing mention).

Citing-text excerpts — how the field used this work

METHODOLOGY Mapping new realities: Ground truth image creation with pix2pix image-to-image translation

“The random jittering process involves resizing the images to 286x286 pixels and cropping them to the final size of 256x256 pixels [10], [11].”

METHODOLOGY Segment anything for videos: A systematic survey

“Research on SAM mainly adapts the following routes: model compression [28], ensuring model robustness [29], advancing efficient finetuning techniques [4], and developing innovative applications [9], [30] (as illustrated in Fig.)”

■ FOLLOW-UP WORK

Shield: An evaluation benchmark for face spoofing and forgery detection with multimodal large language models

2025 · Visual Intelligence 3 (1), 9, 2025 · 69 citations (GS)

Field-normalised: 44 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	Poison as cure: Visual noise for mitigating object hallucinations in lvms	Alibaba Group, Westlake University, Xiamen University	China	—
2	Can chatgpt detect deepfakes? a study of using multimodal large language models for media forensics	The Chinese University of Hong Kong, Shenzhen, University at Albany, State University of New York, University at Buffalo, State University of New York	China, United States	Background
3	Can we leave deepfake data behind in training deepfake detector?	Tencent, The Chinese University of Hong Kong, Shenzhen, Wuhan University	China	—
4	A survey on evaluation of multimodal large language models	Nanyang Technological University	Singapore	Background
5	X2-dfd: A framework for explainable and extendable deepfake detection	The Chinese University of Hong Kong, Shenzhen, University at Buffalo, State University of New York	China, United States	—
6	Guard me if you know me: Protecting specific face-identity from deepfakes	Peking University, Shenzhen University, Tencent	China	—
7	Ffaa: Multimodal large language model based explainable open-world face forgery analysis assistant	The Chinese University of Hong Kong, Tsinghua University	China	—
8	Forensics-bench: A comprehensive forgery detection benchmark suite for large vision language models	Alibaba, MEGVII Technology, Shanghai AI Laboratory	China, Hong Kong	—
9	Large multimodal models evaluation: a survey	Shanghai AI Lab, Shanghai Artificial Intelligence Laboratory	China	—
10	Facexbench: Evaluating multimodal llms on face understanding	Johns Hopkins University	United States	—
11	A survey on multimodal benchmarks: In the era of large ai models	Hong Kong University of Science and Technology, Zhejiang University	China, Hong Kong	Background
12	Facellm: A multimodal large language model for face understanding	Idiap Research Institute	Switzerland	—
13	Interpretable face anti-spoofing: Enhancing generalization with multimodal large language models	Baidu Inc, Baidu Inc., Chinese Academy of Sciences	China	—
14	Mllm-enhanced face forgery detection: A vision-language fusion solution	Chinese Academy of Sciences, CMCC, UCAS	China	—

No.	Citing paper	Citing institution(s)	Country	S2
15	Multimodal large language models to support real-world fact-checking	Mohamed bin Zayed University of Artificial Intelligence	United Arab Emirates	—
16	Survey on ai-generated media detection: From non-mlm to mllm	Beijing University of Posts and Telecommunications, Chinese Academy of Sciences, ShanghaiTech University	China, United States	Influential
17	Large multimodal agents for accurate phishing detection with enhanced token optimization and cost reduction	American University of Beirut	Lebanon	Influential
18	Exploring chatgpt for face presentation attack detection in zero and few-shot in-context learning	Idiap Research Institute	Switzerland	—
19	IKA2: Internal Knowledge Adaptive Activation for Robust Recognition in Complex Scenarios	Beijing Jiaotong University, Jinggangshan University, Tianjin University of Science and Technology	China	—
20	Forenx: Towards explainable ai-generated image detection with multimodal large language models	Beijing Jiaotong University, Microsoft Research Asia	China	—
21	Evolving from single-modal to multi-modal facial deepfake detection: Progress and challenges	Agency for Science, Technology and Research, University of Nevada, Reno	Singapore, United States	—
22	DevFD: Developmental Face Forgery Detection by Learning Shared and Orthogonal LoRA Subspaces	China Mobile Financial Technology Co., Ltd., UCAS, University of Chinese Academy of Sciences	China	—
23	Benchmarking multimodal large language models for face recognition	Idiap Research Institute	Switzerland	—
24	Fit for Purpose? Deepfake Detection in the Real World	Purdue University	United States	—
25	ResProto-FD: Visual-Language Residual Prototype Sets for Generalized Face Forgery Detection	Xidian University	China	—
26	Spark: Multi-vision sensor perception and reasoning benchmark for large-scale vision-language models	KAIST, NVIDIA	South Korea, United States	—
27	Learning to Tell Apart: Weakly Supervised Video Anomaly Detection via Disentangled Semantic Alignment	Huazhong University of Science and Technology	China	—
28	Evaluating Multimodal Large Language Models for Heterogeneous Face Recognition	Idiap Research Institute	Switzerland	—
29	Demographic Fairness in Multimodal LLMs: A Benchmark of Gender and Ethnicity Bias in Face Verification	Idiap Research Institute	Switzerland	—
30	In-context Learning of Vision Language Models for Detection of Physical and Digital Attacks against Face Recognition Systems	Technical University of Denmark, Technische Universität Darmstadt	Denmark, Germany	—

Showing the 30 most-cited of 41 independent citing papers.

Independent citing papers only; self- and co-author citations excluded. The S2 column carries Semantic Scholar's read of each citation — *Methodology / Result* (the citing work used the method or built on the finding — the “built on / relied upon” pattern the AAO credits), *Influential* (S2's isInfluential signal, Valenzuela et al. 2015), or *Background* (a passing mention).

FOLLOW-UP WORK

GM-DF: Generalized multi-scenario deepfake detection

2025 · Proceedings of the 33rd ACM International Conference on Multimedia, 4300-4309, 2025 · 38 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	Evolving from single-modal to multi-modal facial deepfake detection: Progress and challenges	Agency for Science, Technology and Research, University of Nevada, Reno	Singapore, United States	—
2	Unmasking synthetic realities in generative AI: A comprehensive review of adversarially robust deepfake detection systems	Hamad Bin Khalifa University	Qatar	—
3	Evaluating modern approaches in 3d scene reconstruction: Nerf vs gaussian-based methods	Columbia University, Cornell University, Guangxi University	China, Germany, United States	Background
4	Mindev: Multi-modal integrated diffusion framework for video reconstruction from eeg signals	Bytedance Inc., Shanghai Jiao Tong University, University of Shanghai for Science and Technology	China	—
5	Optimizing Management Strategies for Enhanced Performance and Energy Efficiency in Modern Computing Systems	NXP Semiconductor	Netherlands	—
6	The impact of network load balancing on organizational efficiency and managerial decision-making in digital enterprises	Meta Platforms	United States	—
7	Towards sar automatic target recognition: Multi-category sar image classification based on light weight vision transformer	Khalifa University, Swinburne University of Technology, University of Canterbury	Australia, China, New Zealand	—
8	Open-set deepfake detection: a parameter-efficient adaptation method with forgery style mixture	City University of Hong Kong, Hong Kong Baptist University, Jiangxi University of Finance and Economics	Brazil, China, Singapore	—
9	Semantics-oriented multitask learning for DeepFake detection: A joint embedding approach	City University of Hong Kong, Nanyang Technological University, University at Buffalo, State University of New York	China, Singapore, United States	—
10	Deep learning for lung disease classification using transfer learning and a customized cnn architecture with attention	Arizona State University, University of Illinois at Chicago	United States	Methodology
11	Securing social media against deepfakes using identity, behavioral, and geometric signatures	University of Michigan-Flint	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
12	StegaFFD: Privacy-preserving Face Forgery Detection via Fine-grained Steganographic Domain Lifting	Beihang University, Chinese Academy of Sciences, Great Bay University	China, Singapore	—
13	DeepFake detection in images and videos: a survey on models, datasets, and evaluation metrics	Islamic University of Science & Technology	India	—
14	SFANet: Spatial-Frequency Attention Network for Deepfake Detection	Indian Institute of Technology Ropar, Monash University	Australia, India	—
15	Face2Parts: Exploring Coarse-to-Fine Inter-Regional Facial Dependencies for Generalized Deepfake Detection	Korea Aerospace University	South Korea	—
16	Interactive dual-branch network based on adversarial knowledge distillation for compressed deepfake detection: G. Yang et al.	The First Affiliated Hospital, Anhui University of Science and Technology	China	—
17	CLaRE: CLIP with Latent Reconstruction Errors for Generated Face Detection	University of Amsterdam	Netherlands	—
18	From Spurious to Causal: Low-rank Orthogonal Subspace Intervention for Generalizable Face Forgery Detection	Nanjing University of Information Science and Technology	China	—
19	Listening Deepfake Detection: A New Perspective Beyond Speaking-Centric Forgery Analysis	Beijing Institute of Technology, University of Science and Technology Beijing	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).

Citing-text excerpts — how the field used this work

METHODOLOGY Deep learning for lung disease classification using transfer learning and a customized cnn architecture with attention

“For CV, applications like real-time pill identification [10], snore sound analysis [11], and deepfake detection [12], [13] are good examples.”

Contribution 2

Claim — Contribution 2

The researcher developed a fully open framework for multimodal training, significantly advancing the democratization of access to large-scale vision-language model development.

The researcher's primary contribution is the development of a fully open framework for multimodal training, as detailed in the 2025 paper 'Llava-onevision-1.5: Fully open framework for democratized multimodal training.' This work stands as a seminal piece in the field, establishing a foundational approach to making advanced multimodal AI training accessible and transparent.

This line of work appears to address the critical need for open-source alternatives in multimodal AI, a domain often dominated by proprietary systems. By providing a fully open framework, the researcher likely lowered barriers to entry for other scientists and developers, fostering a more inclusive and collaborative research environment. The title suggests a focus on democratization, implying that the framework enables broader participation in training and experimenting with vision-language models.

The significance of this contribution is underscored by its rapid uptake within the academic community. With 121 citations, the paper has clearly influenced subsequent research. Notably, 100% of the 333 citing papers classified for this scholar are from independent researchers, indicating that the work has resonated widely beyond the researcher's immediate circle. This high

level of independent citation demonstrates that the framework has become a standard reference or tool for diverse groups of scientists, validating its broad impact and utility in advancing the field of multimodal AI.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 101 · 10 flagged influential by Semantic Scholar

CORE PAPER

[Llava-onevision-1.5: Fully open framework for democratized multimodal training](#)

2025 · arXiv preprint arXiv:2509.23661, 2025 · 121 citations (GS)

Field-normalised: 121 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2025 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	A survey on agentic multimodal large language models	Chinese University of Hong Kong, City University of Hong Kong, Communication University of China	China, Singapore	—
2	Tinted Frames: Question Framing Blinds Vision-Language Models	University of British Columbia, University of California, Berkeley	Canada, United States	—
3	DFLOP: A Data-driven Framework for Multimodal LLM Training Pipeline Optimization	Microsoft, SK Telecom, Yonsei University	South Korea, United States	—
4	ViCToR: Improving Visual Comprehension via Token Reconstruction for Pretraining LMMs	Imperial College London	United Kingdom	—
5	Revisiting Multimodal KV Cache Compression: A Frequency-Domain-Guided Outlier-KV-Aware Approach	Fudan University, The Chinese University of Hong Kong, Zhangjiang Laboratory	China	Influential
6	Are VLMs Lost Between Sky and Space? LinkSBench for UAV-Satellite Dynamic Cross-View Spatial Intelligence	Qinghai Normal University, Sun Yat-sen University, Xidian University	China	—
7	Multi-Crit: Benchmarking Multimodal Judges on Pluralistic Criteria-Following	University of Maryland, College Park, University of Waterloo	Canada, United States	—
8	T2i-reasonbench: Benchmarking reasoning-informed text-to-image generation	The Chinese University of Hong Kong, The University of Hong Kong	China, Hong Kong	—
9	Think with 3d: Geometric imagination grounded spatial reasoning from limited views	Beihang University, Meituan, National University of Singapore	China, Singapore	—
10	Vismem: Latent vision memory unlocks potential of vision-language models	Fudan University, National University of Singapore, Tsinghua University	China, Singapore	Influential
11	Openmmreasoner: Pushing the frontiers for multimodal reasoning with an open and general recipe	MiroMind AI, Nanyang Technological University, Shanda Group	China, Singapore	—
12	Towards cross-view point correspondence in vision-language models	Beihang University, Beijing Academy of Artificial Intelligence, Chinese Academy of Sciences	China, Singapore	—

No.	Citing paper	Citing institution(s)	Country	S2
13	Deepgen 1.0: A lightweight unified multimodal model for advancing image generation and editing	Fudan University, Nanjing University, Shanghai Innovation Institute	China	—
14	Llada-o: An effective and length-adaptive omni diffusion model	Ant Group, Renmin University of China	China	Influential
15	Thinking in dynamics: How multimodal large language models perceive, track, and reason dynamics in physical 4d world	CUHK, Shenzhen University, SUSTech	China, United States	—
16	InfiniteVL: Synergizing Linear and Sparse Attention for Highly-Efficient, Unlimited-Input Vision-Language Models	Horizon Robotics, Huazhong University of Science and Technology	China	—
17	Audio-Visual Intelligence in Large Foundation Models	Microsoft Research, National University of Singapore, Queen Mary University of London	Canada, China, Singapore	—
18	Small drafts, big verdict: Information-intensive visual reasoning via speculation	New York University, University of California, San Diego	United States	Influential
19	VINO: A Unified Visual Generator with Interleaved OmniModal Context	Kuaishou Technology, Nanyang Technology University, Shanghai AI Laboratory	China, Singapore	—
20	SlowBA: An efficiency backdoor attack towards VLM-based GUI agents	Shanghai Jiao Tong University	China	—
21	Scaling Agentic Reinforcement Learning for Tool-Integrated Reasoning in VLMs	Cisco, Eigen AI, Emory University	Saudi Arabia, United States	—
22	Token Reduction via Local and Global Contexts Optimization for Efficient Video Large Language Models	Tsinghua University, University of Electronic Science and Technology of China, University of Trento	China, Italy	—
23	What if? Emulative simulation with world models for situated reasoning	Hunan University, Karlsruhe Institute of Technology	China, Germany	—
24	Robo-Dopamine: General Process Reward Modeling for High-Precision Robotic Manipulation	Beijing Academy of Artificial Intelligence, Beijing Academy of Artificial Intelligence; University of Sydney, Chinese Academy of Sciences	China, China; Australia	—
25	LogicOCR: Do Your Large Multimodal Models Excel at Logical Reasoning on Text-Rich Images?	Beihang University, Nanyang Technological University, Wuhan University	China, Singapore	—
26	One Model, Two Minds: Task-Conditioned Reasoning for Unified Image Quality and Aesthetic Assessment	Jiigan Technology, Monash University, UESTC	Australia, China	—
27	A Simple Baseline for Unifying Understanding, Generation, and Editing via Vanilla Next-token Prediction	Meituan, Peking University, University of Chinese Academy of Sciences	China	—
28	Remember Me: Bridging the Long-Range Gap in LVLMs with Three-Step Inference-Only Decay Resilience Strategies	Beijing Normal Kong Baptist University, Shanghai Jiao	Australia, China	—

No.	Citing paper	Citing institution(s)	Country	S2
		Tong University, University of Wollongong		
29	NeSy-Route: A Neuro-Symbolic Benchmark for Constrained Route Planning in Remote Sensing	Nanjing University	China	—
30	Tuna-2: Pixel Embeddings Beat Vision Encoders for Multimodal Understanding and Generation	Meta AI, The University of Hong Kong, University of Waterloo	Canada, Hong Kong, United States	—

Showing the 30 most-cited of 101 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 established a comprehensive framework for full-stack LLM agent safety, addressing critical gaps across data, training, and deployment phases.

The researcher's core contribution is anchored in the 2025 paper 'A comprehensive survey in llm (-agent) full stack safety: Data, training and deployment.' This work appears to synthesize the fragmented landscape of large language model safety into a unified, end-to-end perspective. By explicitly covering data, training, and deployment, the titles indicate a holistic approach that addresses the entire lifecycle of LLM agents, rather than isolated components.

This line of work appears to address a significant gap in the literature by providing a structured overview of safety mechanisms across the full stack. The absence of follow-up papers by the same researcher suggests this survey serves as a definitive, standalone reference point that consolidated existing knowledge at a critical juncture in the field's development.

The significance of this contribution is evidenced by its rapid uptake, with 136 citations recorded. Notably, 100% of the 333 classified citing papers originate from independent researchers, indicating that the work has been widely adopted and utilized by the broader scientific community outside the researcher's immediate circle. This high degree of independent citation underscores the paper's role as a foundational resource for diverse teams working on LLM safety.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 92 · 1 flagged influential by Semantic Scholar

CORE PAPER

[A comprehensive survey in llm \(-agent\) full stack safety: Data, training and deployment](#)

2025 · arXiv preprint arXiv:2504.15585, 2025 · 136 citations (GS)

Field-normalised: 113 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2025 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	A survey of llm-driven ai agent communication: Protocols, security risks, and defense countermeasures	Ant Group, ByteDance, Changsha University of Science and Technology	China, Saudi Arabia, United States	—
2	Diagnosing Knowledge Conflict in Multimodal Long-Chain Reasoning	Beijing University of Posts and Telecommunications, Huazhong University	China, Japan, Singapore	—

No.	Citing paper	Citing institution(s)	Country	S2
		of Science and Technology, Nanyang Technological University		
3	Defense-to-attack: Bypassing weak defenses enables stronger jailbreaks in Vision-Language Models	City University of Hong Kong, Fudan University	China	—
4	Trism for agentic ai: A review of trust, risk, and security management in llm-based agentic multi-agent systems	Cornell University, University of Groningen, Vector Institute	Canada, Netherlands, United States	—
5	Towards reasoning era: A survey of long chain-of-thought for reasoning large language models	Central South University, Fudan University, Harbin Institute of Technology	China	—
6	Alphaedit: Null-space constrained knowledge editing for language models	National University of Singapore, University of Science and Technology of China, Zhejiang University	China, Singapore	—
7	Agentauditor: Human-level safety and security evaluation for llm agents	KTH Royal Institute of Technology, Nanyang Technological University, National University of Singapore	Australia, Singapore, Sweden	—
8	Unique security and privacy threats of large language models: A comprehensive survey	City University of Macau, CSIRO, University of Illinois at Chicago	Australia, China, United States	—
9	Advancing expert specialization for better moe	Beijing University of Posts and Telecommunications, Nanyang Technological University	China, Singapore	—
10	Vibe coding vs. agentic coding: Fundamentals and practical implications of agentic ai	Cornell University, University of the Peloponnese	Greece, United States	—
11	Graph-augmented large language model agents: Current progress and future prospects	Griffith University, Nanyang Technological University, National University of Singapore	Australia, China, Singapore	—
12	Red queen: Safeguarding large language models against concealed multi-turn jailbreaking	Hippocratic AI, University of California San Diego, University of Southern California	United States	—
13	From prompt injections to protocol exploits: Threats in LLM-powered AI agents workflows	De Montfort University, Guelma University, Khalifa University of Science and Technology	Algeria, United Arab Emirates, United Kingdom	Influential
14	Cchall: A novel benchmark for joint cross-lingual and cross-modal hallucinations detection in large language models	Central South University, Fudan University, National University of Singapore	China, Singapore	—
15	On the resilience of llm-based multi-agent collaboration with faulty agents	Carnegie Mellon University, Chinese University of Hong Kong, City University of Hong Kong, Shenzhen	China, United States	—
16	Lifebench: Evaluating length instruction following in large language models	Beijing University of Posts and Telecommunications,	China, Japan, Singapore	—

No.	Citing paper	Citing institution(s)	Country	S2
		North China Electric Power University, NTU		
17	Crabs: Consuming resource via auto-generation for llm-dos attack under black-box settings	Beijing University of Posts and Telecommunications, Nanyang Technological University	China, Singapore	—
18	Security of internet of agents: Attacks and countermeasures	Xi'an Jiaotong University	China	—
19	Evaluating the usability of microgestures for text editing tasks in virtual reality	The Hong Kong University of Science and Technology (Guangzhou), University of Cambridge	China, United Kingdom	—
20	A comprehensive survey on trustworthiness in reasoning with large language models	University of Chinese Academy of Sciences	China	—
21	Llm-based agentic reasoning frameworks: A survey from methods to scenarios	Beijing Jiaotong University, Lancaster University, Max Planck Institute for Informatics	China, Germany, United Kingdom	—
22	The eye of sherlock holmes: Uncovering user private attribute profiling via vision-language model agentic framework	Alibaba Group, Beijing University of Technology, Hengxin Tech	China, Singapore	—
23	We should identify and mitigate third-party safety risks in mcp-powered agent systems	National University of Singapore, Shandong University, Tsinghua University	China, Singapore	—
24	Safesieve: From heuristics to experience in progressive pruning for llm-based multi-agent communication	Chongqing University, Nanyang Technological University, National University of Singapore	China, Singapore	—
25	Get experience from practice: Llm agents with record & replay	Shanghai Jiao Tong University	China	—
26	Collaborative inference and learning between edge slms and cloud LLMs: A survey of algorithms, execution, and open challenges	Hong Kong University of Science and Technology, Huazhong University of Science and Technology, Nanyang Technological University	China, Hong Kong, Singapore	—
27	Three minds, one legend: Jailbreak large reasoning model with adaptive stacked ciphers	MIT, Nanyang Technological University, Tsinghua University	China, Singapore, United States	—
28	Malsight: Exploring malicious source code and benign pseudocode for iterative binary malware summarization	Beijing University of Posts and Telecommunications	China	—
29	Quality-diversity red-teaming: Automated generation of high-quality and diverse attackers for large language models	Hong Kong University of Science and Technology, Nanjing University, Southern University of Science and Technology	China, Hong Kong	—

No.	Citing paper	Citing institution(s)	Country	S2
30	Goal-aware identification and rectification of misinformation in multi-agent systems	Beijing University of Posts and Telecommunications, National University of Singapore, University of Science and Technology of China	China, Singapore	—

Showing the 30 most-cited of 92 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
Zhejiang University	China	SCImago #6 · THE 39 · QS 49	36
Nanyang Technological University	Singapore	SCImago #137	34
National University of Singapore	Singapore	SCImago #59 · THE 17 · QS 8	29
Shanghai Jiao Tong University	China	SCImago #10 · THE 40 · QS =47	23
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	22
University of Science and Technology of China	China	SCImago #77 · THE 51 · QS =132	22
Peking University	China	SCImago #11 · THE 13 · QS 14	18
Fudan University	China	SCImago #46 · THE 36 · QS 30	17
Chinese Academy of Sciences	China	SCImago #2	15
Beijing University of Posts and Telecommunications	China	SCImago #355 · QS 1001-1200	14
Sun Yat-sen University	China	SCImago #40 · THE 201–250 · QS =276	14
Nanjing University	China	SCImago #178 · THE =62 · QS =103	11
Beihang University	China	SCImago #160 · THE 251–300 · QS =388	10
The Chinese University of Hong Kong	Hong Kong	SCImago #163 · THE =41 · QS =32	10
East China Normal University	China	SCImago #769 · THE 251–300 · QS =433	9

Geographic distribution of citing authors

Country	Citing papers
China	229
United States	87
Singapore	59

Country	Citing papers
Hong Kong	18
Australia	15
United Kingdom	13
India	12
United Arab Emirates	11
South Korea	8
Germany	8
Japan	7
Switzerland	7

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	Detect any deepfakes: Segment anything meets face forgery detection and localization	91	Dhanasar – Prong 2 (well-positioned)
Contribution 2	Llava-onevision-1.5: Fully open framework for democratized multimodal training	101	Dhanasar – Prong 2 (well-positioned)
Contribution 3	A comprehensive survey in llm (-agent) full stack safety: Data, training and deployment	92	Dhanasar – Prong 2 (well-positioned)