

Citation Evidence Report

EB-2 NIW Petition — National Interest Waiver

Matter of Dhanasar · Prong 2 (well-positioned)

Junxian Li

Shanghai Jiao Tong 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

174 Citing papers mapped	175 Citation edges	23 Home papers mapped	7 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.

70.1% independent of 174 classified citing papers

Citation type	Count
Independent	122
Self-citation	8
Co-author	36
Same-institution	8

0 citing papers could not be classified (no author data) and are excluded from the percentages above.

C. Significant Contributions & Their Citation Evidence

Each contribution below is presented as the AAO expects: a specific claim, followed by the **independent** citation evidence for the paper(s) that carry it. Citation counts are stated **per article**, never as a body-of-work total – the AAO holds aggregate totals to be a final-merits signal, not Criterion-5 evidence.

Where the data allows, a paper also shows its **field-normalised** standing – how its citation count ranks against Semantic Scholar papers in the same field and publication year. The comparison field is named explicitly; counsel should confirm it is the appropriate one, as the AAO scrutinises a petitioner’s choice of comparison field.

Contribution 1

Claim – Contribution 1

The researcher pioneered multimodal large language models for chemistry and subsequently investigated efficiency vulnerabilities in VLM-based GUI agents, establishing a foundational line of inquiry in specialized AI applications.

The researcher's contribution centers on the development and critical analysis of multimodal large language models (VLMs) in specialized domains. This line of work is anchored by the core paper, "Chemvml: Exploring the power of multimodal large language models in chemistry area" (2025), which appears to establish a framework for applying VLMs to chemical sciences. The work was extended by a follow-up study, "SlowBA: An efficiency backdoor attack towards VLM-based GUI agents" (2026), suggesting a broader investigation into the robustness and security of such models in interactive environments.

This trajectory indicates an original approach to bridging general-purpose multimodal AI with domain-specific applications like chemistry, while simultaneously addressing emerging security concerns in VLM-based agents. The progression from capability exploration in chemistry to efficiency attacks in GUI agents suggests a comprehensive examination of both the utility and vulnerabilities of these advanced models.

The significance of this work is evidenced by its reception in the academic community. The core paper has accumulated 90 citations, indicating strong engagement with the proposed methodology. Furthermore, citation analysis reveals that 75.3% of the 174 classified citations originate from independent researchers, demonstrating that the work has influenced scholars outside the researcher's immediate institution and collaboration network, thereby confirming its independent impact and broad relevance.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 53 · 3 flagged influential by Semantic Scholar

CORE PAPER

[Chemvml: Exploring the power of multimodal large language models in chemistry area](#)

2025 · 90 citations (GS)

Field-normalised: 68 Semantic Scholar citations place it in the top 5% of Chemistry papers from 2025 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	From text to insight: large language models for chemical data extraction (2024)	Friedrich Schiller University	Germany	—
2	MicroVQA: A Multimodal Reasoning Benchmark for Microscopy-Based Scientific Research (2025)	Chan Zuckerberg Biohub Network, KTH Royal Institute of Technology, Princeton University	China, Sweden, United States	—
3	MP-GUI: Modality Perception with MLLMs for GUI Understanding (2025)	Ant Group, Zhejiang University	China	—
4	Leveraging Biomolecule and Natural Language through Multi-Modal Learning: A Survey (2024)	Huazhong University of Science and Technology, Renmin University of China, Shanghai AI Laboratory	China	—
5	Multimodal Large Language Models Meet Multimodal Emotion Recognition and Reasoning: A Survey (2025)	Central South University of Forestry and Technology, State University of New York at New Paltz	China, United States	—
6	Exploiting Edited Large Language Models as General Scientific Optimizers (2025)	University of Science and Technology of China	China	—

No.	Citing paper	Citing institution(s)	Country	S2
7	Bee: A High-Quality Corpus and Full-Stack Suite to Unlock Advanced Fully Open MLLMs (2025)	Beihang University, Tencent, Tsinghua University	China	—
8	A survey on large language models in biology and chemistry (2025)	Seoul National University	South Korea	—
9	TinyChemVL: Advancing Chemical Vision-Language Models via Efficient Visual Token Reduction and Complex Reaction Tasks (2026)	Institute of Automation, Chinese Academy of Sciences, University of Chinese Academy of Sciences	China	Influential
10	Efficiently Integrate Large Language Models with visual perception: A survey from the training paradigm perspective (2026)	Lingnan University	Hong Kong	—
11	LLMs4All: A Review of Large Language Models Across Academic Disciplines (2025)	University of Notre Dame	United States	—
12	CultureCLIP: Empowering CLIP with Cultural Awareness through Synthetic Images and Contextualized Captions (2025)	Hong Kong University of Science and Technology, University of Copenhagen, University of Massachusetts Amherst	Denmark, Hong Kong, United States	—
13	LatentChem: From Textual CoT to Latent Thinking in Chemical Reasoning (2026)	—	—	—
14	Large Language Models Meet Molecules: A Systematic Review of Advances and Challenges in AI-Driven Cheminformatics (2025)	National University of Sciences and Technology (NUST), NovaMechanics Ltd, University of Birmingham	Cyprus, Pakistan, United Kingdom	—
15	Use of Machine Learning and Large Language Models in Chemical Information Extraction. (2026)	Hong Kong University of Science and Technology	China	—
16	ChemAU: Harness the Reasoning of LLMs in Chemical Research with Adaptive Uncertainty Estimation (2025)	Fudan University, The Chinese University of Hong Kong	China	—
17	Cognitive Mismatch in Multimodal Large Language Models for Discrete Symbol Understanding (2026)	Guangdong Laboratory of AI and Digital Economy, Independent Researcher, Sun Yat-sen University	China, United States	—
18	Stress-Testing Multimodal Foundation Models for Crystallographic Reasoning (2025)	Hamad Bin Khalifa University, Texas A&M University, Texas A&M University at Qatar	Qatar, United States	—
19	MAPS: Multi-Agent Personality Shaping for Collaborative Reasoning (2026)	Nanyang Technological University, National University of Singapore, Xi'an Jiaotong University	China, Singapore	—
20	Foundation Models for Scientific Discovery: From Paradigm Enhancement to Paradigm Transition (2025)	The Hong Kong University of Science and Technology, The Hong Kong University of Science and Technology (Guangzhou)	China	—

No.	Citing paper	Citing institution(s)	Country	S2
21	Can Vision Language Models Assess Graphic Design Aesthetics? A Benchmark, Evaluation, and Dataset Perspective (2026)	—	—	—
22	MarsRetrieval: Benchmarking Vision-Language Models for Planetary-Scale Geospatial Retrieval on Mars (2026)	Southern University of Science and Technology	China	—
23	Sustainable Materials Discovery in the Era of Artificial Intelligence (2026)	Imperial College London, Indian Institute of Technology Delhi	India, United Kingdom	—
24	New Opportunities For the Integration of Artificial Intelligence With Materials Science: From Large Language Models to Embodied Large Models (2026)	Chemical Defense Institute, National University of Defense Technology	China	—
25	Hierarchical Process Reward Models are Symbolic Vision Learners (2025)	Adelaide AIML, Data61 CSIRO, NetMind.ai	Australia, United Kingdom, United States	—
26	MultiTox: a multimodal fusion model for molecular toxicity prediction. (2025)	Sungkyunkwan University	South Korea	—
27	ChemVLR: Prioritizing Reasoning in Perception for Chemical Vision-Language Understanding (2026)	Chinese Academy of Sciences and University of Chinese Academy of Sciences, Institute of Automation, Chinese Academy of Sciences	China	—
28	ChemVTS-Bench: Evaluating Visual-Textual-Symbolic Reasoning of Multimodal Large Language Models in Chemistry (2025)	Beijing University of Posts and Telecommunications, Gaotu Techedu Inc, Renmin University of China	China	Influential
29	Emergent Pose-Invariance in 3D Molecular Representations via Multimodal Learning (2025)	IBM Research	Japan	—
30	El Agente Estructural: An Artificially Intelligent Molecular Editor (2026)	Acceleration Consortium, University of Toronto	Canada	—

Showing the 30 most-cited of 52 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

[SlowBA: An efficiency backdoor attack towards VLM-based GUI agents](#)

2026 · 1 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	PEARL: Personalized Streaming Video Understanding Model (2026)	—	—	—

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 reinforcement learning methods for traffic signal control that remain effective despite missing data, addressing a critical robustness gap in intelligent transportation systems.

The researcher's core contribution centers on the 2023 paper titled 'Reinforcement learning approaches for traffic signal control under missing data.' This work stands as the primary artifact in this specific line of inquiry, with no subsequent follow-up papers by the same author listed in the provided data. The title suggests a focus on enhancing the reliability of automated traffic management systems when sensor data is incomplete or unreliable.

This line of work appears to address the practical challenge of data scarcity in real-world traffic environments. By applying reinforcement learning to scenarios with missing data, the researcher likely introduced methods that improve system resilience compared to traditional approaches that assume complete information. The absence of follow-up papers indicates this contribution is a distinct, self-contained advancement in the field.

The significance of this work is evidenced by its citation record. With 32 citations, the paper has attracted attention from the academic community. Notably, 75.3% of the 174 citing papers classified for this scholar originate from independent researchers, suggesting that the methodology or findings have been adopted and built upon by peers outside the researcher's immediate institution or collaboration network.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 15 · 1 flagged influential by Semantic Scholar

CORE PAPER

[Reinforcement learning approaches for traffic signal control under missing data](#)

2023 · 32 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Optimizing Urban Road Networks: A Systematic Review of Design, Control and Multimodal Integration (2025)	Obafemi Awolowo University	Nigeria	—
2	DiffLight: A Partial Rewards Conditioned Diffusion Model for Traffic Signal Control with Missing Data (2024)	—	—	Methodology
3	DiffusionLight: a multi-agent reinforcement learning approach for traffic signal control based on shortcut-diffusion model (2025)	Guangxi University of Science and Technology	China	—
4	RobustLight: Improving Robustness via Diffusion Reinforcement Learning for Traffic Signal Control (2025)	Beijing University of Posts and Telecommunications, Lanzhou University, University of Technology Sydney	Australia, China	—
5	FuzzyLight: A Robust Two-Stage Fuzzy Approach for Traffic Signal Control Works in Real Cities (2025)	Beijing University of Posts and Telecommunications, Griffith University, Lanzhou University	Australia, China	—
6	City-scale high-resolution traffic datasets with refined networks for hierarchical traffic control (2026)	Sun Yat-sen University, Tongji University	China	—
7	Result constraint behavior cloning for offline reinforcement learning (2025)	Jilin University, Northeast Normal University	China	—

No.	Citing paper	Citing institution(s)	Country	S2
8	HALO: Hierarchical Reinforcement Learning for Large-Scale Adaptive Traffic Signal Control (2026)	University of Exeter, University of Warwick	United Kingdom	—
9	VLMLight: Safety-Critical Traffic Signal Control via Vision-Language Meta-Control and Dual-Branch Reasoning Architecture (2025)	Fourier Intelligence, Nanyang Technological University, Nokia Bell Labs	China, France, Singapore	—
10	Reinforcement Learning for Digital Twins (2024)	Karlsruhe Institute of Technology, Technical University of Denmark, University of Rostock	Denmark, Germany	—
11	An Interdisciplinary and Cross-Task Review on Missing Data Imputation (2025)	The Chinese University of Hong Kong, Shenzhen	China	—
12	“Less Knowledge is Less” - An Empirical Study of Intelligent Traffic Signal Network Efficiency Under Partial Information (2025)	Duke Kunshan University, University of Ottawa, Wuhan University	Canada, China	—
13	Rethinking the Design of Human-Centric AI Systems for Deployment in Transportation (2025)	Carnegie Mellon University, The University of Texas at Austin, University of Pittsburgh	United States	—
14	A fuzzy-based planning method and meta-heuristic optimization for traffic signal control (2025)	Islamic Azad Iran University, Islamic Azad University, Shahed University	Iran	—
15	DQN (2025)	Qingdao University of Technology	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 DiffLight: A Partial Rewards Conditioned Diffusion Model for Traffic Signal Control with Missing Data

“Similar to the existing work in [14], we adopt store-and-forward method (SFM) [42], a rule-based method that has generally more stable performances, to impute missing observations and rewards for these approaches.”

Contribution 3

Claim — Contribution 3

The researcher developed Critic-V, a framework using VLM critics to detect errors in multimodal reasoning, establishing a novel approach to improving model reliability.

The researcher’s contribution centers on the development of Critic-V, introduced in a 2025 paper titled ‘Critic-v: Vlm critics help catch vlm errors in multimodal reasoning.’ This work proposes a mechanism where vision-language model critics are employed to identify and correct reasoning errors within multimodal systems. The titles indicate a focus on enhancing the robustness and accuracy of these complex AI models through self-critique or external critique mechanisms.

This line of work appears to address the challenge of error detection in multimodal reasoning, a critical gap in current VLM capabilities. By introducing a critic-based approach, the researcher offers a distinct methodological perspective on validating

model outputs. The absence of follow-up papers in the provided data suggests this core publication stands as the primary articulation of this specific contribution, highlighting its foundational nature in the researcher's portfolio.

The significance of this work is evidenced by its citation record, with 48 citations recorded for the core paper. Furthermore, analysis of the broader citation context reveals that 75.3% of citing papers originate from independent researchers, indicating that the methodology has been adopted and built upon by the wider scientific community beyond the researcher's immediate circle. This high degree of independent uptake underscores the utility and impact of the Critic-V framework in advancing multimodal AI research.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 30 · 2 flagged influential by Semantic Scholar

CORE PAPER

Critic-v: Vlm critics help catch vlm errors in multimodal reasoning

2025 · 48 citations (GS)

Field-normalised: 42 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	InternLM-XComposer2.5-Reward: A Simple Yet Effective Multi-Modal Reward Model (2025)	Shanghai Artificial Intelligence Laboratory	China	—
2	PRMBench: A Fine-grained and Challenging Benchmark for Process-Level Reward Models (2025)	—	—	—
3	MM-Verify: Enhancing Multimodal Reasoning with Chain-of-Thought Verification (2025)	—	—	—
4	UniCTokens: Boosting Personalized Understanding and Generation via Unified Concept Tokens (2025)	—	—	—
5	Explain Before You Answer: A Survey on Compositional Visual Reasoning (2025)	Griffith University, Monash University, Princeton University	Australia, United States	—
6	Generative Universal Verifier as Multimodal Meta-Reasoner (2025)	ByteDance, Princeton University, Tsinghua University	China, United States	—
7	Sherlock: Self-Correcting Reasoning in Vision-Language Models (2025)	Purdue University	United States	—
8	Large-scale multimodal model based embodied intelligent robots: A survey (2026)	Fudan University, Hunan University, Northeastern University	China, United States	—
9	ViPER: Empowering the Self-Evolution of Visual Perception Abilities in Vision-Language Models (2025)	MBZUAI, Meituan, Renmin University of China	China, United Arab Emirates	—
10	Insight-V++: Towards Advanced Long-Chain Visual Reasoning with Multimodal Large Language Models (2026)	Nanyang Technological University, Tencent Hunyuan, Tsinghua University	China, Singapore	—
11	MARBLE: A Hard Benchmark for Multimodal Spatial Reasoning and Planning (2025)	EPFL	Switzerland	—

No.	Citing paper	Citing institution(s)	Country	S2
12	OCR-Critic: Aligning Multimodal Large Language Models' Perception through Critical Feedback (2025)	Beijing University of Posts and Telecommunications, Renmin University of China, South China University of Technology	China	—
13	Make LVLMS Focus: Context-Aware Attention Modulation for Better Multimodal In-Context Learning (2026)	Brown University, Massachusetts Institute of Technology, University of Bristol	United Kingdom, United States	—
14	MMC: Iterative Refinement of VLM Reasoning via MCTS-based Multimodal Critique (2025)	iFLYTEK Research, University of Science and Technology of China	China	—
15	They Said Memes Were Harmless --- We Found the Ones That Hurt: Decoding Jokes, Symbols, and Cultural References (2026)	Jamia Hamdard, Macquarie University, Stanford University	Australia, India, United States	Influential
16	Agentic World Modeling: Foundations, Capabilities, Laws, and Beyond (2026)	—	—	—
17	GAIA: A Data Flywheel System for Training GUI Test-Time Scaling Critic Models (2026)	—	—	—
18	Understanding the Implicit User Intention via Reasoning with Large Language Model for Image Editing (2025)	Johns Hopkins University, Southern University of Science and Technology	China, United States	—
19	Cross-Domain Demo-to-Code via Neurosymbolic Counterfactual Reasoning (2026)	Sungkyunkwan University	—	—
20	Clarification as Supervision: Reinforcement Learning for Vision-Language Interfaces (2025)	—	—	—
21	Counteracting Matthew Effect in Self-Improvement of LVLMS through Head-Tail Rebalancing (2025)	—	—	—
22	Toward Artificial Metacognition (2026)	Syracuse University	United States	—
23	CPR-CLIP: Cross-modal Consistent and Prompt-diverse Regularized CLIP for Action Recognition (2025)	Wuhan Textile University	China	—
24	Cognitive Hives: A Distributed Systems Theory for Conflict-Aware Multimodal Intelligence (2026)	Independent Researcher	United States	Influential
25	Look Before You Leap: A GUI-Critic-R1 Model for Pre-Operative Error Diagnosis in GUI Automation (2025)	—	—	—
26	GENIUS: Generative Fluid Intelligence Evaluation Suite (2026)	—	—	—
27	HermesFlow: Seamlessly Closing the Gap in Multimodal Understanding and Generation (2025)	—	—	—
28	Learning to Pose Problems: Reasoning-Driven and Solver-Adaptive Data Synthesis for Large Reasoning Models (2025)	Tencent, Tsinghua University	China	—

No.	Citing paper	Citing institution(s)	Country	S2
29	From Perception to Reasoning: Deep Thinking Empowers Multimodal Large Language Models (2025)	Harbin Institute of Technology	China	—
30	When Meaning Isn't Literal: Exploring Idiomatic Meaning Across Languages and Modalities (2026)	—	—	—

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
Shanghai Jiao Tong University	China	SCImago #10 · THE 40 · QS =47	25
Fudan University	China	SCImago #46 · THE 36 · QS 30	14
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	14
Shanghai AI Laboratory	China	—	11
Shanghai Artificial Intelligence Laboratory	China	SCImago #563	8
Nanyang Technological University	Singapore	SCImago #137	7
Arizona State University	United States	SCImago #357 · THE 201–250 · QS =173	7
The Hong Kong Polytechnic University	Hong Kong	SCImago #256 · THE 80 · QS 54	7
National University of Singapore	Singapore	SCImago #59 · THE 17 · QS 8	6
Zhejiang University	China	SCImago #6 · THE 39 · QS 49	6
University of Science and Technology of China	China	SCImago #77 · THE 51 · QS =132	6
Nanjing University	China	SCImago #178 · THE =62 · QS =103	5
Sun Yat-sen University	China	SCImago #40 · THE 201–250 · QS =276	5
Institute of Automation, Chinese Academy of Sciences	China	SCImago #340	5
The Chinese University of Hong Kong	China	SCImago #163 · THE =41 · QS =32	5

Geographic distribution of citing authors

Country	Citing papers
China	96
United States	39
Hong Kong	14
Singapore	11

Country	Citing papers
Australia	10
United Kingdom	7
Germany	4
Canada	3
Japan	3
Switzerland	2
United Arab Emirates	2
Denmark	2

Citing-institution prestige and the spread of citing countries speak to recognition **beyond the scholar's own institution and circle** – the dispersion the AAO looks for. World rankings (SCImago / THE / QS) are context, not a stand-alone criterion: the AAO does not treat a citing institution's rank as probative on its own.

E. Citation Growth Over Time

Distinct citing papers by publication year. Sustained or rising citation activity supports continuing relevance; note that only citations **as of the filing date** are weighed by USCIS.



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	Chemvlm: Exploring the power of multimodal large language models in chemistry area	53	Dhanasar — Prong 2 (well-positioned)
Contribution 2	Reinforcement learning approaches for traffic signal control under missing data	15	Dhanasar — Prong 2 (well-positioned)
Contribution 3	Critic-v: Vlm critics help catch vlm errors in multimodal reasoning	30	Dhanasar — Prong 2 (well-positioned)