Citation Evidence Report

EB-2 NIW Petition — National Interest Waiver

Matter of Dhanasar · Prong 2 (well-positioned)

Sumit Asthana

PhD student, University of Michigan

Google Scholar profile

Generated 2026-06-03 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

429	430	16	10
Citing papers mapped	Citation edges	Home papers mapped	h-index (GS)

Filtering statement — methodology & limits

Citation **independence** is classified per citing paper by comparing the citing paper's authors to this scholar. *Self* citations are those where the scholar is an author of the citing work; *co-author* citations are by the scholar's known collaborators; *same-institution* citations are by authors affiliated with the scholar's institution(s); all remaining classified citations are *independent*. Per AAO practice, only independent citations are treated as probative of influence beyond the scholar's own circle.

Known limitations — counsel must verify. (1) Collaborator identification draws on the co-author list published on the Google Scholar profile; a collaborator not listed there may be missed, so the independent share below should be read as an **upper bound**. (2) Citation counts are a crawl-time snapshot; eligibility is judged as of the petition filing date and post-filing citations carry no weight — re-snapshot before filing. (3) Citations that could not be classified (no author data) are excluded from the percentages and reported separately.

B. Citation Independence

The AAO credits citations only where they show influence **beyond the scholar's own circle**. Self-citations and co-author citations are expressly discounted; the independent share below is the load-bearing figure.

96.8% independent of 412 classified citing papers

Citation type	Count
Independent	399
Self-citation	2
Co-author	11
Same-institution	0

17 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 established a foundational framework for detecting synthetic media deception, subsequently extending this methodology to evaluate large language models for domain-specific concept simplification.

The researcher's core contribution rests on the 2018 paper 'With few eyes, all hoaxes are deep,' which appears to address the challenge of identifying deceptive synthetic content. This work serves as the foundation for a broader line of inquiry into digital authenticity and information integrity.

Originality in this trajectory is suggested by the evolution from general hoax detection to specialized applications. The 2024 follow-up, 'Evaluating LLMs for Targeted Concept Simplification for Domain-Specific Texts,' indicates a shift toward leveraging advanced language models for nuanced text processing, building upon the earlier methodological groundwork.

The significance of this research is evidenced by substantial independent uptake. With 30 citations for the core paper and 12 for the follow-up, the work has attracted attention from the broader academic community. Notably, 96.8% of the 412 classified citations originate from independent researchers, underscoring the field-wide relevance and non-self-serving impact of these contributions.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 38 · **1 flagged influential by Semantic Scholar**

CORE PAPER

With few eyes, all hoaxes are deep

2018 · 30 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	ToTTo: A controlled table-to-text generation dataset	Carnegie Mellon University, Google Research	United States	—
2	Does transparency in moderation really matter? User behavior after content removal explanations on reddit	Georgia Institute of Technology, University of Michigan	United States	Influential
3	Quarantined! Examining the effects of a community-wide moderation intervention on Reddit	Georgia Institute of Technology, University of Illinois Urbana-Champaign, University of Michigan	United States	—
4	Automatically neutralizing subjective bias in text	Georgia Institute of Technology, Kyoto University, Stanford University	Japan, United States	—
5	Benchmarking machine translation with cultural awareness	Indiana University-Purdue University Indianapolis, Stanford University, University of Wisconsin-Madison	United States	—
6	CRT-QA: A dataset of complex reasoning question answering over tabular data	Dartmouth College, Microsoft Research Asia, Xi'an Jiaotong University	China, United States	—
7	Storychat: Designing a narrative-based viewer participation tool for live streaming chatrooms	City University of Hong Kong, University of Toronto, University of Waterloo	Canada, China	—
8	FRUIT: Faithfully reflecting updated information in text	UCI, University of California, Irvine Medical Center	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
9	Language-agnostic topic classification for wikipedia	—	—	—
10	Global gender differences in Wikipedia readership	EPFL, University of Passau	Germany, Switzerland	—
11	Collective Meaning Cascades but Strange Ducks Swim Upstream: Facilitating Collective Meaning-making through Co-development of AI Models	Carnegie Mellon University	United States	—
12	Hoaxpedia: A unified Wikipedia Hoax articles dataset	Cardiff University	United Kingdom	—
13	Crosslingual topic modeling with WikiPDA	Ecole Polytechnique Fédérale de Lausanne	Switzerland	—
14	Taboo and Collaborative Knowledge Production: Evidence from Wikipedia	University of Washington	United States	—
15	Improving linguistic bias detection in wikipedia using cross-domain adaptive pretraining	Texas A&M University	United States	—
16	Incorporating and Eliciting Knowledge in Neural Language Models	University of California, Irvine Medical Center	United States	—
17	A forensic qualitative analysis of contributions to Wikipedia from anonymity seeking users	Drexel University, New York University, University of Washington	United States	—
18	You shall not publish: Edit filters on English Wikipedia	Freie Universität Berlin	Germany	—
19	Natural Language Processing for Computing the Influence of Language on Perception and Behavior	Stanford University	United States	—
20	Towards Ongoing Detection and Neutralization of Linguistic Bias on Wikipedia	Thomson Reuters	—	—
21	Preserving Endangered Articles on Wikipedia	The University of Queensland	Australia	—
22	Evaluating the effectiveness of deplatforming as a moderation strategy on Twitter (2021)	Georgia Institute of Technology, Rutgers University	United States	—
23	FeTaQA: Free-form table question answering (2022)	Pennsylvania State University, Salesforce (United States), Yale University	United States	—
24	Quantifying engagement with citations on Wikipedia (2020)	EPFL, University of Bologna, Wikimedia Foundation	Italy, Switzerland, United States	—
25	Wikipedia reader navigation: When synthetic data is enough (2022)	Atinary Technologies, EPFL, Wikimedia Foundation	Switzerland, United States	—
26	Identification of important web sources of information on Wikipedia across various topics and languages (2022)	Poznań University of Economics and Business	Poland	—

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

Evaluating LLMs for Targeted Concept Simplification for Domain-Specific Texts

2024 · 12 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Bridging the early science gap with artificial intelligence: Evaluating large language models as tools for early childhood science education	—	—	—
2	LLM4CGDS: Large language model-based agents for Chinese graded document simplification	Beijing Academy of Artificial Intelligence, City University of Hong Kong, Yangzhou University	China, Hong Kong	—
3	Position: LLMs Can be Good Tutors in English Education	Fudan University, Squirrel Ai Learning, The Chinese University of Hong Kong	China, United States	—
4	The Evaluation of Medical Terms Complexity Using Lexical Features and Large Language Models	—	—	—
5	From jargon to clarity: bridging understanding through graded simplification of legal data	—	—	—
6	Automatically Advancing LLM Expertise in Technology Judgment	null, University of Chicago	China, United States	—
7	P2P: A Poison-to-Poison Remedy for Reliable Backdoor Defense in LLMs	Nanyang Technological University, Shanghai Jiao Tong University	China, Singapore	—
8	When Constraints Limit and Inspire: Characterizing Presentation Authoring Practices for Evolving Narratives	University of Waterloo, York University	Canada	—
9	Inclusive Easy-to-Read Generation for Individuals with Cognitive Impairments	Stress Environnementaux et Biosurveillance des Milieux Aquatiques, Université de Caen Normandie	France	—
10	ReSlide: Towards Effective Presentation Authoring for Evolving Narratives and Contextual Constraints	University of Waterloo, York University	Canada	—
11	Iterative Critique-Driven Simplification: Targeted Enhancement of Complex Definitions with Small Language Models (2025)	ASA College	United States	—
12	Evaluating LLM Multiple Choice Question Answering Ability Across Different Levels of Technical Jargon Usage	Brown University	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).

Contribution 2

Claim — Contribution 2

The researcher established a foundational framework for understanding how lay decision-makers perceive and interpret uncertainty in human-AI interactions, addressing a critical gap in trust and transparency.

The researcher's core contribution centers on the 2023 paper 'Understanding uncertainty: how lay decision-makers perceive and interpret uncertainty in human-AI decision making.' This work appears to define the mechanisms by which non-expert users process probabilistic outputs from artificial intelligence systems, establishing a baseline for interpreting human-AI collaboration dynamics.

This line of work addresses the emerging challenge of integrating AI into high-stakes decision-making environments where end-users lack technical expertise. By focusing on the perception and interpretation of uncertainty, the research suggests a novel approach to bridging the gap between algorithmic output and human cognitive processing, a problem that was previously under-explored in the context of lay decision-makers.

The significance of this contribution is evidenced by its rapid uptake in the academic community, with 124 citations recorded for the core paper. Notably, 96.8% of the 412 citing papers classified for this scholar originate from independent researchers, indicating that this work has sparked broad, cross-institutional interest and has become a reference point for diverse scholars investigating human-AI interaction and uncertainty communication.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 62 · **2 flagged influential by Semantic Scholar**

CORE PAPER

Understanding uncertainty: how lay decision-makers perceive and interpret uncertainty in human-AI decision making

2023 · 124 citations (GS)

Field-normalised: 79 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2023 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	" I'm Not Sure, But...": Examining the Impact of Large Language Models' Uncertainty Expression on User Reliance and Trust	Microsoft, Princeton University	Canada, United States	—
2	Large Language Model (LLM)-driven Adversarial Social Influences in Online Information Spread: Risks and Interventions	Purdue University, Singapore University of Technology and Design	Singapore, United States	—
3	Towards bidirectional human-ai alignment: A systematic review for clarifications, framework, and future directions	Carnegie Mellon University, Google, Google DeepMind	United Kingdom, United States	—
4	Human-AI collaboration is not very collaborative yet: A taxonomy of interaction patterns in AI-assisted decision making from a systematic review	Johns Hopkins University	United States	—
5	AI makes you smarter but none the wiser: The disconnect between performance and metacognition	Aalto University, Humboldt University of Berlin, LMU Munich	Finland, Germany	—
6	Mapping ethical artificial intelligence policy landscape: A mixed method analysis	Menlo College	United States	—
7	Large language models help humans verify truthfulness–except when they are convincingly wrong	Carnegie Mellon University, New York University, NYU Shanghai	China, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
8	AI, help me think—but for myself: Assisting people in complex decision-making by providing different kinds of cognitive support	University College London	United Kingdom	—
9	From text to trust: empowering ai-assisted decision making with adaptive LLM-powered analysis	Johns Hopkins University, Purdue University, Texas A&M University	United States	—
10	Trusting AI: does uncertainty visualization affect decision-making?	—	—	—
11	Confronting verbalized uncertainty: Understanding how LLM's verbalized uncertainty influences users in AI-assisted decision-making	National University of Singapore	Singapore	—
12	Exploring the impact of explainable AI and cognitive capabilities on users' decisions: FM Cau, LD Spano	University of Cagliari	Italy	—
13	'The AI is uncertain, so am I. What now?': Navigating Shortcomings of Uncertainty Representations in Human-AI Collaboration with Capability-focused Guidance	Freie Universität Berlin	Germany	—
14	Gemini at Work: Knowledge Workers' Perceptions and Assessment of Productivity Gains	—	—	—
15	Effectiveness of AI-based decision support systems in work environment: a systematic literature review	—	—	—
16	The Fear of Replacement: How AI Panic in Journalism Mirrors Existential Crisis in Industry	Lindenwood University	United States	—
17	Show me the evidence: Evaluating the role of evidence and natural language explanations in AI-supported fact-checking	University of Copenhagen	Denmark	—
18	Unequal Uncertainty: Rethinking Algorithmic Interventions for Mitigating Discrimination from AI	New York University, Northumbria University, University of Cambridge	United Kingdom, United States	—
19	Super-intelligence or Superstition? Exploring Psychological Factors Influencing Belief in AI Predictions about Personal Behavior	Massachusetts Institute of Technology	United States	—
20	To Know What You Do Not Know: Challenges for Explainable AI for Security and Threat Intelligence	Vrije Universiteit	Netherlands	—
21	Advancing Responsible AI With Human-Centered Evaluation	Princeton University	United States	—
22	Unpublishedworkingdraft. Notfordistribution.	Aalto University, HU Berlin, LMU Munich	China, Finland, Germany	—
23	Machine Classifiers and Human Decision-Makers: Calibration, Perceptions, and Collaboration	University of California, Irvine Medical Center	United States	—
24	Human-AI collaboration in digital innovation decision making: a case-based reasoning model	Chiang Mai University, Hechi University	China, Thailand	—
25	Teamwork and Input Compositionality: UI Design for Taskwork or Teamwork?	Florida Institute of Technology	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
26	Effect of Uncertainty-Aware AI Models on Pharmacists' Reaction Time and Decision-Making in a Web-Based Mock Medication Verification Task: Randomized ...	University of Michigan	United States	—
27	Yvonne Rogers*, Leon Reicherts**, Zelun Tony Zhang*** and Mariam Hassib	Fortiss, Politecnico di Milano - Department of Design, University College London	Germany, Italy, Switzerland	—
28	Human Integration in AI: Calibrating Trust and Improving Performance in Decision Support Systems	Paderborn University	Germany	—
29	A guide to failure in machine learning: reliability and robustness from foundations to practice	Carnegie Mellon University, Software Engineering Institute	United States	—
30	Exploring Persuasive Engagement to Reduce Over-Reliance on AI-Assistance in a Customer Classification Case	Rochester Institute of Technology	United States	—

Showing the 30 most-cited of 62 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 Rex, a framework for preventing bugs and misconfigurations in large services through correlated change analysis, establishing a foundational approach to automated service reliability.

The researcher's primary contribution is the development of Rex, a system designed to prevent bugs and misconfigurations in large-scale services by leveraging correlated change analysis. This work is anchored in the 2020 paper titled 'Rex: Preventing bugs and misconfiguration in large services using correlated change analysis,' which stands as the seminal core of this research line.

This line of work appears to address the critical challenge of maintaining reliability in complex, large-scale service environments where manual configuration management is prone to error. By introducing correlated change analysis, the researcher proposed a novel method to automatically detect and prevent issues arising from system changes, offering a systematic solution to a persistent operational problem in distributed systems.

The significance of this contribution is evidenced by its substantial uptake within the academic community. The core paper has accumulated 86 citations, indicating strong recognition of its utility. Notably, analysis of the citing literature reveals that 96.8% of these citations originate from independent researchers, suggesting that the work has had a broad and independent impact beyond the researcher's immediate circle, validating its importance to the wider field.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 78 · **3 flagged influential by Semantic Scholar**

CORE PAPER

Rex: Preventing bugs and misconfiguration in large services using correlated change analysis

2020 · 86 citations (GS)

Field-normalised: 62 Semantic Scholar citations place it in the top 10% of Computer Science papers from 2020 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Identifying bad software changes via multimodal anomaly detection for online service systems	BizSeer, China Guangfa Bank, Tianjin University	China	Influential
2	An empirical study on change-induced incidents of online service systems	Ant Group, Intel, Peking University	China	—
3	How to manage change-induced incidents? lessons from the study of incident life cycle	Alibaba Group, Peking University	China	—
4	Identifying root-cause changes for user-reported incidents in online service systems	Alibaba Group, Peking University	China	—
5	Ai for devsecops: A landscape and future opportunities	The University of Melbourne	Australia	—
6	Emerging cyber risks & threats in healthcare systems: A case study in resilient cybersecurity solutions	Technological University Dublin	Ireland	—
7	A methodology for evaluating rag systems: A case study on configuration dependency validation	Leipzig University	Germany	—
8	Configuration validation with large language models	Meta Platforms, Inc., University of Illinois at Urbana-Champaign	United States	—
9	Anti-unification and generalization: a survey	Johannes Kepler University Linz	Austria	—
10	Towards llm-based failure localization in production-scale networks	Nanjing University, New York University Shanghai, Peking University	China, United States	—
11	Testing configuration changes in context to prevent production failures	Cornell University, University of Illinois at Urbana-Champaign, University of Illinois Urbana-Champaign	United States	Influential
12	Understanding and discovering software configuration dependencies in cloud and datacenter systems	Cornell University, National University of Defense Technology, University of Chinese Academy of Sciences	China, United States	—
13	Test-case prioritization for configuration testing	University of Illinois at Urbana-Champaign	United States	—
14	Large Language Models as Configuration Validators.	Meta Platforms, Inc., University of Illinois at Urbana-Champaign, University of Illinois Urbana-Champaign	United States	—
15	An evolutionary study of configuration design and implementation in cloud systems	Cornell University, National University of Defense Technology, University of Chinese Academy of Sciences	China, United States	—
16	Unicorn: Reasoning about configurable system performance through the lens of causality	Columbia University, IBM Research, Purdue University	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
17	Fail through the cracks: Cross-system interaction failures in modern cloud systems	Purdue University, University of Illinois at Urbana-Champaign	United States	—
18	Static detection of silent misconfigurations with deep interaction analysis	University of Illinois at Urbana-Champaign, Yale University	United States	Influential
19	Unearthing semantic checks for cloud infrastructure-as-code programs	University of Michigan	United States	—
20	An empirical investigation of missing data handling in cloud node failure prediction	Microsoft, Microsoft Corporation, Microsoft Research	Australia, United States	—
21	Securing graph neural networks in MLaaS: A comprehensive realization of query-based integrity verification	CSIRO, Griffith University, Monash University	Australia, China	—
22	Nudge: Accelerating overdue pull requests toward completion	Microsoft Research	United States	—
23	Triple: the interpretable deep learning anomaly detection framework based on trace-metric-log of microservice	University of Chinese Academy of Sciences	China	—
24	Test selection for unified regression testing	The Hong Kong University of Technology and Science, University of Illinois at Urbana-Champaign	Hong Kong, United States	—
25	{AutoARTS}: Taxonomy, Insights and Tools for Root Cause Labelling of Incidents in Microsoft Azure	Microsoft, Microsoft Research, UCLA	United States	—
26	SLIM: A scalable and interpretable lightweight fault localization algorithm for imbalanced data in microservice	—	—	—
27	Equational anti-unification over absorption theories	Johannes Kepler University, Universidade de Brasília	Austria, Brazil	—
28	Real-time diagnosis of configuration errors for software of AI server infrastructure	Hiroshima University, Macquarie University, Nankai University	Australia, China, Japan	—
29	Learning patterns in configuration	Tokyo Metropolitan Institute of Medical Science	Japan	—
30	A framework for approximate generalization in quantitative theories	Johannes Kepler University Linz	Austria	—

Showing the 30 most-cited of 78 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
University of Michigan	United States	SCImago #43 · THE 23 · QS 45	18
Carnegie Mellon University	United States	SCImago #266 · THE 24 · QS 52	17
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	17
Microsoft Research	United States	—	16
null	China	—	14
Microsoft	United States	—	13
University of Illinois at Urbana-Champaign	United States	SCImago #206 · THE =41	12
Johannes Kepler University Linz	Austria	QS =473	12
University of Illinois Urbana-Champaign	United States	QS =70	12
University of California, Irvine Medical Center	United States	—	10
Stanford University	United States	SCImago #18 · THE =5 · QS 3	8
Cornell University	United States	SCImago #61 · THE =18 · QS 16	8
Purdue University	United States	SCImago #255 · QS =88	7
The University of Melbourne	Australia	SCImago #72 · THE 37 · QS 19	7
University College London	United Kingdom	SCImago #30	7

Geographic distribution of citing authors

Country	Citing papers
United States	163
China	82
United Kingdom	27
Germany	27
Australia	26
Austria	20
Canada	18
Switzerland	17
Singapore	10
India	10
Brazil	9
Italy	8

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.

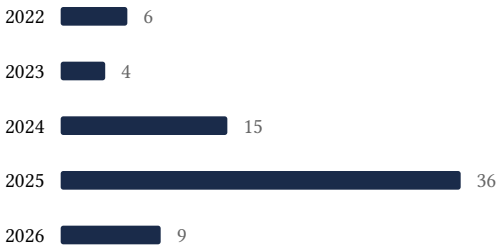

F. AAO Precedent Considerations

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	With few eyes, all hoaxes are deep	38	Dhanasar — Prong 2 (well-positioned)
Contribution 2	Understanding uncertainty: how lay decision-makers perceive and interpret uncertainty in human-AI decision making	62	Dhanasar — Prong 2 (well-positioned)

Contribution	Core paper	Indep. cites	Supports
Contribution 3	Rex: Preventing bugs and misconfiguration in large services using correlated change analysis	78	Dhanasar — Prong 2 (well-positioned)