Citation Evidence Report

EB-2 NIW Petition — National Interest Waiver

Matter of Dhanasar · Prong 2 (well-positioned)

Jeffrey Wu

Anthropic AI, OpenAI

Google Scholar profile

Generated 2026-06-08 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

627	627	31	17
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.

98.2% independent of 614 classified citing papers

Citation type	Count
Independent	603
Self-citation	3
Co-author	8
Same-institution	0

13 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 scalable few-shot learning in language models and advanced instruction-following capabilities through human feedback, establishing foundational methods for generalizing model behavior.

CLAIM: The researcher's contribution centers on advancing the capabilities of large language models, anchored by the seminal 2020 paper "Language models are few-shot learners." This work appears to have established a critical baseline for how models can perform tasks with minimal examples, a theme extended in subsequent publications.

ORIGINALITY: The chronological progression from few-shot learning to instruction following with human feedback suggests a strategic evolution in model alignment. The titles indicate a shift from passive pattern recognition to active, human-guided behavior modification. The later work on weak-to-strong generalization further implies an effort to elicit robust capabilities using less intensive supervision, addressing the challenge of scaling model utility efficiently.

SIGNIFICANCE: The impact of this line of work is evidenced by the core paper's 76,558 citations and the follow-up's 26,749 citations, indicating widespread adoption. With 98.2% of classified citations originating from independent researchers, the work demonstrates broad, field-wide influence beyond the researcher's immediate circle, confirming its status as a foundational contribution to the field.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 603 · **105 flagged influential by Semantic Scholar**

CORE PAPER

Language models are few-shot learners

2020 · 76,558 citations (GS)

Field-normalised: 57,162 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2020 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	The rise and potential of large language model based agents: A survey	Fudan University, Meituan	China	Influential
2	Accurate predictions on small data with a tabular foundation model	University of Freiburg	Germany	—
3	Tree of thoughts: Deliberate problem solving with large language models	Google DeepMind, Princeton University, University of California, Irvine Medical Center	United Kingdom, United States	—
4	Alpacafarm: A simulation framework for methods that learn from human feedback	Cornell University, Stanford, Stanford University	Canada, United States	—
5	Diffusion models: A comprehensive survey of methods and applications	Beijing University of Posts and Telecommunications, Carnegie Mellon University, Peking University	China, United States	—
6	Mamba: Linear-time sequence modeling with selective state spaces	Princeton University	United States	—
7	Ai alignment: A comprehensive survey	Peking University	China	—
8	A comprehensive survey of small language models in the era of large language models: Techniques, enhancements, applications, collaboration with llms, and ...	Amazon, Penn State University, The Pennsylvania State University	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
9	Attention mechanisms in computer vision: A survey	Cardiff University, Nankai University, Tsinghua University	China	—
10	ChatGPT utility in healthcare education, research, and practice: systematic review on the promising perspectives and valid concerns	University of Jordan	Jordan	—
11	Mm-llms: Recent advances in multimodal large language models	Kyoto University, Mohamed bin Zayed University of Artificial Intelligence, Tencent	China, Japan, United Arab Emirates	—
12	Voxposer: Composable 3d value maps for robotic manipulation with language models	MIT, Stanford University, University of Pennsylvania	United States	—
13	Modelscope text-to-video technical report	Alibaba Group, Huazhong University of Science and Technology	China	—
14	Yolov12: Attention-centric real-time object detectors	University at Buffalo, University of Chinese Academy of Sciences	China, United States	—
15	Instruction tuning for large language models: A survey	Alibaba Group, Nanyang Technological University, Peking University	China, Singapore, United States	Influential
16	A comprehensive survey of loss functions and metrics in deep learning	Centro de Investigaciones en Óptica A.C., Instituto Politecnico Nacional, Universidad Autónoma de Querétaro	Mexico	—
17	Federated learning for generalization, robustness, fairness: A survey and benchmark	Hong Kong University of Science and Technology, Wuhan University	China, Hong Kong	—
18	Visual prompt tuning	Cornell University, Google DeepMind, Meta	United States	Influential
19	Visual cot: Advancing multi-modal language models with a comprehensive dataset and benchmark for chain-of-thought reasoning	SenseTime, Sensetime (China), SenseTime Research	Canada, China, Hong Kong	—
20	How to build the virtual cell with artificial intelligence: Priorities and opportunities	Broad Institute, Chan Zuckerberg Initiative (United States), Enzo Life Sciences	Canada, France, Germany	—
21	Hyenadna: Long-range genomic sequence modeling at single nucleotide resolution	Harvard University, Mila and Université de Montréal, Stanford University	Canada, United States	Influential
22	Nucleotide transformer: building and evaluating robust foundation models for human genomics	InstaDeep	United Kingdom	—
23	Transformers and genome language models	Stanford University, University of Toronto	Canada, United States	—
24	Empowering biomedical discovery with AI agents	Broad Institute, Harvard Medical School, Harvard University	United Kingdom, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
25	Sequential modeling enables scalable learning for large vision models	Indian Institute of Technology Kanpur, Johns Hopkins University, UC Berkeley	India, Israel, United States	—
26	AI Art and its Impact on Artists	Artist, California State University Northridge, Distributed Artificial Intelligence Research Institute	Canada, United States	—
27	Towards a general-purpose foundation model for computational pathology	Brigham and Women's Hospital, Harvard Medical School	United States	—
28	The current and future state of AI interpretation of medical images	Harvard Medical School	United States	—
29	Vision-language models for medical report generation and visual question answering: A review	H. Lee Moffitt Cancer Center and Research Institute	United States	—
30	From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair NLP models	Carnegie Mellon University, University of Washington, Xi'an Jiaotong University	China, United States	**Influential**

Showing the 30 most-cited of 603 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

Training language models to follow instructions with human feedback

2022 · 26,749 citations (GS)

Field-normalised: 20,077 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2022 indexed by Semantic Scholar, by citation count.

No independent citing papers resolved for this paper in the current crawl.

FOLLOW-UP WORK

Weak-to-strong generalization: Eliciting strong capabilities with weak supervision

2023 · 528 citations (GS)

Field-normalised: 437 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2023 indexed by Semantic Scholar, by citation count.

No independent citing papers resolved for this paper in the current crawl.

Contribution 2

Claim — Contribution 2

The researcher advanced the field by demonstrating that large-scale language models can perform complex tasks with minimal examples, establishing few-shot learning as a viable paradigm for general-purpose AI.

The researcher's core contribution rests on the seminal 2020 paper 'Language models are few-shot learners,' published in Advances in Neural Information Processing Systems. This work appears to have fundamentally shifted the understanding of how large-scale models can be utilized, suggesting that extensive task-specific training is not always necessary for high performance.

This line of work addresses the challenge of adapting powerful language models to diverse tasks without costly fine-tuning. By focusing on few-shot capabilities, the research indicates a move toward more flexible and efficient AI systems that can generalize from limited examples, a significant departure from previous methodologies that relied heavily on supervised learning for each specific task.

The significance of this contribution is underscored by its substantial citation count of 77,508, reflecting widespread adoption and influence within the scientific community. Furthermore, analysis of 614 citing papers reveals that 98.2% originate from independent researchers, indicating that the work has driven broad, external innovation and is not merely the result of self-citation or institutional clustering.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 0

CORE PAPER

Language models are few-shot learners

2020 · Advances in neural information processing systems 33, 1877-1901, 2020 · 77,508 citations (GS)

Field-normalised: 57,162 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2020 indexed by Semantic Scholar, by citation count.

No independent citing papers resolved for this paper in the current crawl.

Contribution 3

Claim — Contribution 3

The researcher established that language models function as unsupervised multitask learners, a foundational framework that has profoundly influenced the development of modern natural language processing systems.

The researcher's seminal contribution rests on the 2019 paper titled 'Language models are unsupervised multitask learners.' This work posits that language models can effectively serve as unsupervised multitask learners, offering a unified approach to various natural language processing tasks without requiring task-specific supervised training data.

This line of work appears to address the fragmentation in natural language processing, where distinct models were traditionally required for different tasks. By proposing a single, unsupervised framework capable of handling multiple tasks, the researcher introduced a novel paradigm that challenges the necessity of specialized, supervised architectures for each specific linguistic challenge.

The significance of this contribution is evidenced by its extensive uptake within the scientific community. With over 39,000 citations, the paper has become a cornerstone reference in the field. Furthermore, citation analysis reveals that 98.2% of citing papers originate from independent researchers, indicating that this work has driven broad, cross-institutional innovation and established a new standard for understanding language model capabilities.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 0

CORE PAPER

Language models are unsupervised multitask learners

2019 · 39,151 citations (GS)

Field-normalised: 28,234 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2019 indexed by Semantic Scholar, by citation count.

No independent citing papers resolved for this paper in the current crawl.

D. Citing-Institution Prestige & Geography

Top citing institutions

Institution	Country	World ranking	Citing papers
Stanford University	United States	SCImago #18 · THE =5 · QS 3	53
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	39
University of California, Irvine Medical Center	United States	—	36
Peking University	China	SCImago #11 · THE 13 · QS 14	35
Carnegie Mellon University	United States	SCImago #266 · THE 24 · QS 52	32
Google Research	United States	—	32
UC Berkeley	United States	—	27
National University of Singapore	Singapore	SCImago #59 · THE 17 · QS 8	24
Princeton University	United States	SCImago #386 · THE =3 · QS =25	24
Microsoft	United States	—	22
Google DeepMind	United States	SCImago #90	21
Google	United States	—	20
The Chinese University of Hong Kong	Hong Kong	SCImago #163 · THE =41 · QS =32	19
University of Washington	United States	SCImago #45 · THE 25 · QS 81	18
Tencent	United States	—	18

Geographic distribution of citing authors

Country	Citing papers
United States	362
China	224
United Kingdom	66
Hong Kong	43
Singapore	39
Canada	39
Germany	36
Australia	29
Switzerland	15
France	12
Japan	11
United Arab Emirates	11

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	Language models are few-shot learners	603	Dhanasar — Prong 2 (well-positioned)
Contribution 2	Language models are few-shot learners	0	Dhanasar — Prong 2 (well-positioned)
Contribution 3	Language models are unsupervised multitask learners	0	Dhanasar — Prong 2 (well-positioned)