

Citation Evidence Report

EB-2 NIW Petition — National Interest Waiver

Matter of Dhanasar · Prong 2 (well-positioned)


Joelle Pineau

School of Computer Science, McGill University; FAIR, Meta AI; Mila

[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



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.

96.1% independent of 2,586 classified citing papers

Citation type	Count
Independent	2,485
Self-citation	19
Co-author	82
Same-institution	0

251 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 advanced deep reinforcement learning by establishing foundational principles and addressing critical challenges in sample efficiency and environmental sustainability.

The researcher's contribution centers on a seminal 2018 paper titled 'Deep reinforcement learning that matters,' which serves as the cornerstone for a sustained line of inquiry into the practical and theoretical robustness of deep reinforcement learning systems. This core work appears to have defined key criteria for evaluating the significance and reliability of such models.

Originality in this body of work is suggested by the chronological progression from foundational theory to specific, pressing challenges. The titles indicate a shift toward addressing sample efficiency in model-free reinforcement learning from images in 2021, and a parallel focus on the systematic reporting of energy and carbon footprints in machine learning in 2020. This trajectory suggests an effort to broaden the scope of reinforcement learning research beyond pure performance metrics to include resource efficiency and environmental impact.

The significance of this line of work is evidenced by substantial citation counts, with the core paper accumulating 3,359 citations and follow-up works garnering 1,119 and 617 citations respectively. Furthermore, citation analysis reveals that 96.1% of citations to the researcher's work originate from independent researchers, indicating broad adoption and influence across the global academic community rather than localized or self-referential engagement.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 740 · 31 flagged influential by Semantic Scholar

CORE PAPER

[Deep reinforcement learning that matters](#)

2018 · 3,359 citations (GS)

Field-normalised: 2,275 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2018 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Network planning with deep reinforcement learning	Facebook Inc., Johns Hopkins University, Peking University	China, United States	—
2	Quantifying Language Models' Sensitivity to Spurious Features in Prompt Design or: How I learned to start worrying about prompt formatting	University of California Berkeley, University of Washington	United States	—
3	Safe learning in robotics: From learning-based control to safe reinforcement learning	University of Toronto	Canada	—
4	REFORMS: Consensus-based Recommendations for Machine-learning-based Science	Cornell University, Duke University, Ghent University	Belgium, Norway, United Kingdom	—
5	Learning to maximize mutual information for dynamic feature selection	Nanyang Technological University, University of Washington	Singapore, United States	—
6	Reinforcement learning algorithms: A brief survey	Indian Institute of Technology Roorkee	India	—
7	Concepts of artificial intelligence for computer-assisted drug discovery	ETH Zurich	Switzerland	—

No.	Citing paper	Citing institution(s)	Country	S2
8	An introductory review of deep learning for prediction models with big data	Tampere University, University of Applied Sciences Upper Austria	Austria, Finland	—
9	Reinforcement learning with augmented data	Korea Advanced Institute of Science & Technology, New York University, UC Berkeley	South Korea, United States	—
10	Better than classical? the subtle art of benchmarking quantum machine learning models	Chalmers University of Technology, Xanadu	Sweden	—
11	A review of molecular representation in the age of machine learning	University of Cambridge	United Kingdom	—
12	A survey on deep learning for human activity recognition	McGill University, Shenzhen University, The University of Melbourne	Australia, Canada, China	—
13	Machine learning in real-time Internet of Things (IoT) systems: A survey	Baidu, Baidu (China), Baidu Research	China, United States	—
14	Unsupervised representation learning in deep reinforcement learning: A review	University of Twente	Netherlands	—
15	Bayesian generational population-based training	Amazon, Google DeepMind, University of Oxford	Australia, United Kingdom	—
16	A survey on data-efficient algorithms in big data era	Moulay Ismail University	Morocco	—
17	Applications, promises, and pitfalls of deep learning for fluorescence image reconstruction	Chan Zuckerberg Biohub	United States	—
18	Artificial neural networks for photonic applications—from algorithms to implementation: tutorial	Aston University	United Kingdom	—
19	Autocompress: An automatic dnn structured pruning framework for ultra-high compression rates	Midea Group, Northeastern University, Syracuse University	China, United States	—
20	Deep reinforcement learning	Tsinghua University	China	—
21	Benchmark data repositories for better benchmarking	University of California, Irvine Medical Center	United States	—
22	Foundation models in smart agriculture: Basics, opportunities, and challenges	Jeonbuk National University, Michigan State University, Nanyang Technological University	China, Singapore, South Korea	Influential
23	Hallucination detection in foundation models for decision-making: A flexible definition and review of the state of the art	University of Illinois Urbana-Champaign	United States	—
24	A survey on interpretable reinforcement learning	Huawei, Shanghai Jiao Tong University, Tianjin University	China	—
25	Regularized evolution for image classifier architecture search	Google, Google Research	United States	—
26	What matters in learning from offline human demonstrations for robot manipulation	Georgia Institute of Technology, Stanford University, The University of Texas at Austin	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
27	Reinforcement learning for whole-building HVAC control and demand response	Monash University	Australia	—
28	Evolving deep neural networks using coevolutionary algorithms with multi-population strategy.	Auckland University of Technology	New Zealand	—
29	Dense reward for free in reinforcement learning from human feedback	University of Cambridge, University of Oxford	United Kingdom	—
30	Data-centric ai: Perspectives and challenges	Rice University	United States	—

Showing the 30 most-cited of 740 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).

FOLLOW-UP WORK

[Towards the systematic reporting of the energy and carbon footprints of machine learning](#)

2020 · 1,119 citations (GS)

Field-normalised: 660 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.

FOLLOW-UP WORK

[Improving sample efficiency in model-free reinforcement learning from images](#)

2021 · 617 citations (GS)

Field-normalised: 505 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2021 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 pioneered end-to-end generative hierarchical neural models for dialogue systems, establishing a foundational framework that significantly advanced the field of conversational AI.

The researcher's core contribution rests on the 2016 paper 'Building end-to-end dialogue systems using generative hierarchical neural network models,' which introduced a novel architectural approach to dialogue generation. This work appears to have shifted the paradigm toward more coherent, structured conversational agents by leveraging hierarchical neural networks.

This line of work addresses the challenge of generating contextually appropriate and structurally sound dialogue. The subsequent 2017 paper, 'A hierarchical latent variable encoder-decoder model for generating dialogues,' suggests a refinement of this approach by incorporating latent variables, indicating an iterative effort to improve model expressiveness and control. The 2019 paper, 'The second conversational intelligence challenge (convai2),' further demonstrates the researcher's role in standardizing evaluation metrics and fostering community-wide benchmarking for these systems.

The significance of this contribution is evidenced by the substantial citation counts, with the core paper accumulating 2272 citations and the follow-up work receiving 1388 and 488 citations respectively. Notably, 96.1% of the 2586 classified citations originate from independent researchers, underscoring the broad adoption and impact of this framework across the global academic community beyond the researcher's immediate circle.

CORE PAPER

Building end-to-end dialogue systems using generative hierarchical neural network models

2016 · 2,272 citations (GS)

Field-normalised: 1,806 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2016 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Neural networks and deep learning	Brown University, University of Pennsylvania	United States	—
2	Recent trends in deep learning based natural language processing	Nanyang Technological University, National University of Singapore, Singapore University of Technology and Design	Singapore	—
3	A state-of-the-art survey on deep learning theory and architectures	Comcast, Lawrence Livermore National Laboratory, Saint Louis University	United States	—
4	Survey of the state of the art in natural language generation: Core tasks, applications and evaluation	Tilburg University, University of Malta	Malta, Netherlands	—
5	The history began from alexnet: A comprehensive survey on deep learning approaches	Lawrence Livermore National Laboratory, Saint Louis University, St. Jude Children's Research Hospital	United States	—
6	Comprehensive exploration of synthetic data generation: A survey	Argonne National Laboratory, University of Chicago, University of Ulm	Germany, United States	—
7	Recent advances in deep learning based dialogue systems: A systematic survey	Nanyang Technological University, National University of Singapore	Singapore	Influential
8	Graph neural networks for natural language processing: A survey	Central China Normal University, JD.COM, Jingdong	Canada, China, United States	—
9	Efficiently trainable text-to-speech system based on deep convolutional networks with guided attention	Independent Researcher, National Institute of Informatics, PKSHA Technology, Inc.	United States	—
10	A survey of personalized large language models: Progress and future directions	Huawei, Huawei Technologies Co., Ltd, The Chinese University of Hong Kong	China	—
11	Pre-trained language models and their applications	Baidu, Baidu Inc., Carnegie Mellon University	China, United States	—
12	How do you converse with an analytical chatbot? revisiting gricean maxims for designing analytical conversational behavior	Tableau Research	United States	—
13	Odo: Design of multimodal chatbot for an experiential media system	Arizona State University	United States	—
14	Neural networks based domain name generation	National Institute of Standards and Technology	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
15	Recurrent neural network for predicting transcription factor binding sites	Tongji University	China	—
16	Neural approaches to conversational AI	Google Inc., Microsoft Research	United States	—
17	A survey of evaluation metrics used for nlg systems	Indian Institute of Technology, Madras	India	—
18	Survey on evaluation methods for dialogue systems	UNED, Université Paris-Saclay, University of the Basque Country (UPV/EHU)	France, Spain, Switzerland	Influential
19	Conversational agents: Goals, technologies, vision and challenges	Ariel University, Jerusalem College of Technology	Israel	—
20	Neural path hunter: Reducing hallucination in dialogue systems via path grounding	Northeastern University, University of Alberta	Canada, United States	—
21	Towards knowledge-based recommender dialog system	DAMO Academy, Alibaba Group, Tsinghua University, Zhipu AI	China, United States	—
22	Conversational question answering: A survey	Macquarie University, The University of Adelaide	Australia	—
23	Deep reinforcement learning for sequence-to-sequence models	Virginia Tech, Virginia Tech Services	United States	—
24	Less is more: Learning to refine dialogue history for personalized dialogue generation	Beijing Academy of Artificial Intelligence, Renmin University of China, University of Montreal	Canada, China	—
25	Grounded conversation generation as guided traverses in commonsense knowledge graphs	Northeastern University, Tsinghua University	China, United States	Influential
26	Do massively pretrained language models make better storytellers?	Carnegie Mellon University, Google DeepMind, Stanford University	United Kingdom, United States	—
27	RECAP: Retrieval-enhanced context-aware prefix encoder for personalized dialogue response generation	University of Southern California	United States	—
28	KdConv: A Chinese multi-domain dialogue dataset towards multi-turn knowledge-driven conversation	State Key Laboratory of Intelligent Technology and Systems, Tsinghua University	China, PR China	—
29	Generalization in generation: A closer look at exposure bias	ETH Zürich	Switzerland	—
30	A survey on recent advances in llm-based multi-turn dialogue systems	Sun Yat-sen University	China	—

Showing the 30 most-cited of 787 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

[A hierarchical latent variable encoder-decoder model for generating dialogues](#)

2017 · 1,388 citations (GS)

Field-normalised: 1,138 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2017 indexed by Semantic Scholar, by citation count.

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

FOLLOW-UP WORK

[The second conversational intelligence challenge \(convai2\)](#)

2019 · 488 citations (GS)

Field-normalised: 395 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.

Contribution 3

Claim – Contribution 3

The researcher developed a seminal anytime algorithm for Partially Observable Markov Decision Processes, establishing a foundational point-based approach that significantly advanced computational methods in sequential decision-making under uncertainty.

CLAIM: The researcher's primary contribution is the development of a point-based value iteration algorithm for Partially Observable Markov Decision Processes (POMDPs), as detailed in the 2003 paper titled 'Point-based value iteration: An anytime algorithm for POMDPs.' This work stands as a core, standalone achievement in the field, with no subsequent follow-up papers by the researcher listed in this specific line of inquiry.

ORIGINALITY: The title suggests the work addresses the computational challenges inherent in POMDPs by introducing an 'anytime' algorithmic framework. This implies a novel methodological shift toward point-based approximations, likely offering a more efficient or scalable solution compared to prior exact methods, thereby enabling practical applications in complex decision-making environments where full observability is absent.

SIGNIFICANCE: The impact of this contribution is evidenced by its substantial citation count of 1,616, indicating widespread recognition and utility within the academic community. Furthermore, citation analysis reveals that 96.1% of citing papers originate from independent researchers, demonstrating that the work has served as a critical foundation for diverse, external scholarly efforts rather than merely circulating within the researcher's immediate network.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 0

CORE PAPER

[Point-based value iteration: An anytime algorithm for POMDPs](#)

2003 · 1,616 citations (GS)

Field-normalised: 1,223 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2003 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
University of California, Irvine Medical Center	United States	—	90
Tsinghua University	PR China	SCImago #8 · THE 12 · QS =17	78
Carnegie Mellon University	United States	SCImago #266 · THE 24 · QS 52	77
Peking University	China	SCImago #11 · THE 13 · QS 14	71
Stanford University	United States	SCImago #18 · THE =5 · QS 3	67
Microsoft Research	United States	—	54
Nanyang Technological University	Singapore	SCImago #137	45
McGill University	Canada	SCImago #168 · THE =41 · QS 27	44
Shanghai Jiao Tong University	China	SCImago #10 · THE 40 · QS =47	43
Google	United States	—	42
Harbin Institute of Technology	China	SCImago #56 · THE =131 · QS 256	42
National University of Singapore	Singapore	SCImago #59 · THE 17 · QS 8	36
University of Science and Technology of China	China	SCImago #77 · THE 51 · QS =132	35
Google Research	United States	—	34
Chinese Academy of Sciences	China	SCImago #2	34

Geographic distribution of citing authors

Country	Citing papers
United States	933
China	867
United Kingdom	238
Canada	211
Germany	131
Singapore	121
India	103
Japan	82
Australia	76
France	76
South Korea	72
Italy	71

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	Deep reinforcement learning that matters	740	Dhanasar — Prong 2 (well-positioned)
Contribution 2	Building end-to-end dialogue systems using generative hierarchical neural network models	787	Dhanasar — Prong 2 (well-positioned)
Contribution 3	Point-based value iteration: An anytime algorithm for POMDPs	0	Dhanasar — Prong 2 (well-positioned)