Citation Evidence Report

EB-2 NIW Petition — National Interest Waiver

Matter of Dhanasar · Prong 2 (well-positioned)

neta shaul

Computer Science PhD student, Weizmann Institute of Science

[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

32	32	5	8
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.

93.8% independent of 32 classified citing papers

Citation type	Count
Independent	30
Self-citation	0
Co-author	2
Same-institution	0

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 advanced generative modeling and decision-making through guided flows, a contribution evidenced by a seminal 2023 paper with 97 citations from entirely independent researchers.

The researcher's significant contribution centers on the development of guided flows for generative modeling and decision making, as detailed in their 2023 arXiv preprint. This work stands as a core piece of scholarship in the field, establishing a foundation for integrating flow-based methods with decision-making processes.

This line of work appears to address the challenge of effectively combining generative modeling techniques with decision-making frameworks. By introducing guided flows, the researcher likely provided a novel approach to enhancing the utility of generative models in sequential or decision-oriented tasks, distinguishing this work from standard generative modeling efforts.

The significance of this contribution is underscored by its reception within the academic community. With 97 citations, the paper has attracted substantial attention. Notably, 100% of the classified citing papers originate from independent researchers, indicating that the work has resonated beyond the researcher's immediate circle and has been adopted by the broader scientific community as a valuable resource.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 4

CORE PAPER

Guided Flows for Generative Modeling and Decision Making

2023 · arXiv preprint · 97 citations (GS)

Field-normalised: 90 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	Improving Video Generation with Human Feedback (2025)	Kuaishou Technology, The Chinese University of Hong Kong, Tsinghua University	China	—
2	Offline Model-Based Optimization: Comprehensive Review (2026)	Mila - Quebec AI Institute / University of Montreal, University of Washington	Canada, United States	—
3	Generative Models for Crystalline Materials. (2026)	Karlsruhe Institute of Technology, University of California, Berkeley	Germany, United States	—
4	Affordance-based Robot Manipulation with Flow Matching (2024)	Honda Research Institute EU	—	—

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 introduced Discrete Flow Matching, a novel generative modeling framework published at NeurIPS 2024 that has achieved significant independent scholarly adoption.

The researcher's primary contribution is the development of Discrete Flow Matching, presented in a seminal paper at NeurIPS 2024. This work stands as a foundational piece in the researcher's portfolio, establishing a distinct methodological approach within the field of generative modeling without reliance on subsequent follow-up publications by the same author.

The originality of this contribution appears to lie in its specific formulation of flow matching for discrete data, a problem area that traditionally presents unique challenges compared to continuous domains. By isolating this core paper as the primary vehicle for the contribution, the narrative highlights the self-contained impact and theoretical novelty of the initial proposal, suggesting it addressed a critical gap in existing generative frameworks.

The significance of this work is evidenced by its rapid accumulation of 322 citations shortly after publication. Notably, citation analysis reveals that 100% of the citing papers originate from independent researchers, indicating that the broader scientific community has widely adopted and built upon this framework. This high degree of independent uptake underscores the work's utility and influence beyond the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 8 **· 1 flagged influential by Semantic Scholar**

CORE PAPER

Discrete Flow Matching

2024 · NeurIPS · 322 citations (GS)

Field-normalised: 246 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2024 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Transfusion: Predict the Next Token and Diffuse Images with One Multi-Modal Model (2024)	Meta, Waymo	—	Background
2	Large Language Diffusion Models (2025)	—	—	—
3	Fast-dLLM: Training-free Acceleration of Diffusion LLM by Enabling KV Cache and Parallel Decoding (2025)	NVIDIA	—	—
4	Simple Guidance Mechanisms for Discrete Diffusion Models (2025)	Cornell University	United States	—
5	Scaling up Masked Diffusion Models on Text (2024)	Renmin University of China, Sea AI Lab	China	—
6	LLaDA-V: Large Language Diffusion Models with Visual Instruction Tuning (2025)	Renmin University of China	China	—
7	Remasking Discrete Diffusion Models with Inference-Time Scaling (2025)	—	—	**Influential**
8	Unlocking Guidance for Discrete State-Space Diffusion and Flow Models (2024)	University of California, Berkeley	United States	**Result**

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

RESULT Unlocking Guidance for Discrete State-Space Diffusion and Flow Models

"Finally, DiGress is specific to diffusion and not readily generalizable to flow matching, which has been shown to offer several theoretical and practical advantages over diffusion models (Campbell et al., 2024; Gat et al., 2024; Lipman et al., 2023)."

Contribution 3

Claim — Contribution 3

The researcher provided a foundational guide and code for flow matching, establishing a widely adopted resource that has garnered significant independent scholarly attention.

The researcher's contribution centers on the 2024 arXiv publication titled 'Flow Matching Guide and Code,' which serves as the core work in this line of research. This paper appears to offer practical guidance and implementation resources for flow matching, a technique in generative modeling, thereby addressing the need for accessible tools in this domain.

The originality of this work lies in its synthesis of theoretical concepts with practical code, as suggested by the title. By providing a comprehensive guide, the researcher likely lowered the barrier to entry for other scientists, facilitating the adoption and understanding of flow matching methods without requiring extensive prior implementation experience.

The significance of this contribution is evidenced by its high citation count of 308. Notably, 100% of the classified citing papers originate from independent researchers, indicating that the work has been widely recognized and utilized by the broader scientific community beyond the researcher's immediate circle. This broad independent uptake suggests the guide has become a standard reference in the field.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 9

CORE PAPER

Flow Matching Guide and Code

2024 · arXiv · 308 citations (GS)

Field-normalised: 225 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2024 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	A Survey on Vision-Language-Action Models: An Action Tokenization Perspective (2025)	Peking University	China	—
2	Sundial: A Family of Highly Capable Time Series Foundation Models (2025)	Tsinghua University	China	—
3	Generalizing from SIMPLE to HARD Visual Reasoning: Can We Mitigate Modality Imbalance in VLMs? (2025)	Princeton University	United States	—
4	Diffuse and Disperse: Image Generation with Representation Regularization (2025)	MIT	United States	—
5	Horizon Reduction Makes RL Scalable (2025)	Carnegie Mellon University, Princeton University, University of California, Berkeley	United States	—
6	The Principles of Diffusion Models (2025)	OpenAI, Sony AI, Sony Corporation	United States	—
7	Atom-level enzyme active site scaffolding using RFdiffusion2 (2026)	Massachusetts Institute of Technology, University of Washington	United States	—
8	Steering Your Diffusion Policy with Latent Space Reinforcement Learning (2025)	Amazon, UC Berkeley, University of Washington	United States	—
9	Toward Generalist Neural Motion Planners for Robotic Manipulators: Challenges and Opportunities (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
University of Washington	United States	SCImago #45 · THE 25 · QS 81	3
University of California, Berkeley	United States	SCImago #95 · THE 9 · QS =17	3
Stanford University	United States	SCImago #18 · THE =5 · QS 3	2
Meta	United States	—	2
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	2
Princeton University	United States	SCImago #386 · THE =3 · QS =25	2
Peking University	China	SCImago #11 · THE 13 · QS 14	2
Renmin University of China	China	SCImago #2319	2
Cornell University	United States	SCImago #61 · THE =18 · QS 16	1
Massachusetts Institute of Technology	United States	SCImago #41 · THE 2 · QS 1	1
The Hong Kong Polytechnic University	Hong Kong	SCImago #256 · THE 80 · QS 54	1
MIT	United States	—	1
NVIDIA	United States	—	1
Sea AI Lab	Singapore	—	1
The Chinese University of Hong Kong	China	SCImago #163 · THE =41 · QS =32	1

Geographic distribution of citing authors

Country	Citing papers
United States	12
China	7
Canada	1
Germany	1
Hong Kong	1

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.

2024 4

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	Guided Flows for Generative Modeling and Decision Making	4	Dhanasar — Prong 2 (well-positioned)
Contribution 2	Discrete Flow Matching	8	Dhanasar — Prong 2 (well-positioned)
Contribution 3	Flow Matching Guide and Code	9	Dhanasar — Prong 2 (well-positioned)