Citation Evidence Report

EB-1B Petition — Outstanding Professor or Researcher

8 CFR § 204.5(i)(3) · Authorship + Original Contributions

Oliver Wang

Google DeepMind

Google Scholar profile

Generated 2026-05-21 by CiteMap. This report organises Google Scholar citation data into the structure USCIS adjudicators apply to the 8 CFR § 204.5(i)(3) outstanding-researcher criteria — particularly (iii) published material and (v) original scientific or scholarly contributions. 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

34	**34**	**5**	**61**
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.

82.4% independent of 34 classified citing papers

Citation type	Count
Independent	28
Self-citation	0
Co-author	5
Same-institution	1

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 established deep neural network features as a robust perceptual metric, fundamentally shifting how image similarity is quantified in computer vision.

The researcher's core contribution rests on the 2018 IEEE/CVF CVPR paper, 'The Unreasonable Effectiveness of Deep Features as a Perceptual Metric.' This work appears to propose that features extracted from deep learning models serve as a superior standard for measuring perceptual similarity between images, challenging traditional pixel-based or hand-crafted metrics.

This line of work addresses the gap in accurately aligning computational image differences with human visual perception. By leveraging the representational power of deep networks, the research suggests a novel approach to perceptual loss functions, offering a more nuanced understanding of image quality and similarity than previously available methods.

The significance of this contribution is evidenced by its substantial citation count of over 21,000, indicating widespread adoption across the field. Furthermore, analysis shows that 85.3% of citing papers originate from independent researchers, demonstrating that the work has become a foundational reference point for the broader scientific community rather than just the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 7

CORE PAPER

The Unreasonable Effectiveness of Deep Features as a Perceptual Metric

2018 · 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition · 21,275 citations (GS)

Field-normalised: 17,074 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	None (2025)	Microsoft Research, Tsinghua University, USTC, Microsoft Research	—	—
2	Generative artificial intelligence: a systematic review and applications (2024)	Cardiff Metropolitan University, Delhi Technological University, Delhi Technological University (DTU)	India, United Kingdom	—
3	LGM: Large Multi-view Gaussian Model for High-resolution 3D Content Creation (2024)	Nanyang Technological University, NVIDIA, Peking University	China, Hong Kong, Singapore	**Methodology**
4	Visual Autoregressive Modeling: Scalable Image Generation via Next-Scale Prediction (2024)	ByteDance, Peking University	China	—
5	4D Gaussian Splatting for Real-Time Dynamic Scene Rendering (2024)	Huawei, Huazhong University of Science and Technology	China	Background
6	Animate Anyone: Consistent and Controllable Image-to-Video Synthesis for Character Animation (2024)	Institute for Intelligent Computing, Alibaba Group	—	**Methodology**
7	SuGaR: Surface-Aligned Gaussian Splatting for Efficient 3D Mesh Reconstruction and High-Quality Mesh Rendering (2024)	Ecole des Ponts, Univ Gustave Eiffel, CNRS	France	—

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 LGM: Large Multi-view Gaussian Model for High-resolution 3D Content Creation

> *"VGG-based LPIPS loss [60] to the RGB image: We further apply mean square error loss on the alpha image for faster convergence of the shape:"*

METHODOLOGY Animate Anyone: Consistent and Controllable Image-to-Video Synthesis for Character Animation

> *"For quantitative assessment of image-level quality, SSIM[50], PSNR[18] and LPIPS[62] are employed."*

Contribution 2

Claim — Contribution 2

The researcher advanced multimodal image-to-image translation, establishing a foundational framework for generating diverse outputs from single inputs, as evidenced by a seminal NIPS 2017 paper with over 2,000 citations.

The researcher's primary contribution lies in advancing the field of multimodal image-to-image translation. This work is anchored by a seminal paper published in Advances in Neural Information Processing Systems (NIPS) in 2017, which serves as the cornerstone of this specific line of inquiry. The titles indicate a focus on handling multiple possible outputs for a given input, a complex challenge in generative modeling.

This line of work appears to address the limitation of deterministic models that produce only a single output, thereby introducing a framework for capturing the inherent ambiguity in image translation tasks. By focusing on multimodal aspects, the research suggests a novel approach to generating diverse and plausible results, distinguishing it from prior unimodal methods. The absence of follow-up papers in the provided data highlights the standalone impact and completeness of this initial contribution.

The significance of this work is underscored by its substantial citation count, exceeding 2,000 times, which indicates widespread recognition and utility within the scientific community. Furthermore, analysis of citing papers reveals that approximately 85% of citations originate from independent researchers, rather than the author's immediate collaborators or institution. This high degree of independent uptake demonstrates that the work has become a standard reference point for external scholars, validating its broad influence and originality in the field.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 4 · **1 flagged influential by Semantic Scholar**

CORE PAPER

Toward Multimodal Image-to-Image Translation

2017 · Advances in Neural Information Processing Systems 30 (NIPS 2017) · 2,182 citations (GS)

Field-normalised: 1,441 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2017 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Adding Conditional Control to Text-to-Image Diffusion Models (2023)	Stanford University	United States	—
2	Palette: Image-to-Image Diffusion Models (2022)	Google Research	United States	—
3	StarGAN v2: Diverse Image Synthesis for Multiple Domains (2020)	NAVER Corp.	—	Methodology
4	A Review on Generative Adversarial Networks: Algorithms, Theory, and Applications (2023)	—	—	Background

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** StarGAN v2: Diverse Image Synthesis for Multiple Domains
>
> *"nection between stochastic noise and the generated image for diversity, by marginal matching [1], latent regression [40, 13], and diversity regularization [35, 27]."*

Contribution 3

Claim — Contribution 3

The researcher established a foundational benchmark for detecting CNN-generated images, demonstrating their initial detectability while highlighting the transient nature of such vulnerabilities in computer vision security.

CLAIM: The researcher's core contribution is the identification and analysis of artifacts in CNN-generated images, as presented in the 2020 CVPR paper titled 'CNN-generated images are surprisingly easy to spot... for now.' This work serves as the primary anchor for this line of inquiry, standing alone without direct follow-up publications by the same author in the provided dataset.

ORIGINALITY: The title suggests a critical examination of the reliability of synthetic media detection, addressing the gap between perceived robustness and actual vulnerability in early generative models. By framing the detectability as 'surprisingly easy' yet conditional ('for now'), the work appears to have introduced a nuanced perspective on the evolving arms race between generation and detection techniques, challenging assumptions about the permanence of detection methods.

SIGNIFICANCE: The work has garnered substantial attention, evidenced by 1,840 citations, indicating its status as a highly influential reference in the field. Furthermore, citation analysis reveals that 85.3% of citing papers originate from independent researchers, underscoring the broad adoption of these findings across the global academic community and validating the work's impact beyond the researcher's immediate institutional circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 2

> **CORE PAPER**
>
> **CNN-generated images are surprisingly easy to spot... for now**
>
> 2020 · 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) · 1,840 citations (GS)
>
> Field-normalised: 1,397 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	Media Forensics and DeepFakes: An Overview (2020)	University Federico II of Naples	Italy	Methodology
2	SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations (2021)	Google, Stanford University	United States	Background

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** Media Forensics and DeepFakes: An Overview
>
> *"Instead, in [197] a careful pre- and postprocessing and data augmentation are applied to improve transferability."*

D. Citing-Institution Prestige & Geography

Top citing institutions

Institution	Country	World ranking	Citing papers
Adobe Research	United States	—	3
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	3
Microsoft Research	United States	—	3
Stanford University	United States	SCImago #18 · THE =5 · QS 3	3
Shanghai AI Laboratory	China	—	2
Massachusetts Institute of Technology	United States	SCImago #41 · THE 2 · QS 1	2
DeepSeek-AI	China	—	2
New York University	United States	SCImago #116 · THE =31 · QS 55	2
Carnegie Mellon University	United States	SCImago #266 · THE 24 · QS 52	2
Microsoft	United States	—	2
MIT	United States	—	2
NVIDIA	United States	—	2
Peking University	China	SCImago #11 · THE 13 · QS 14	2
University of Science and Technology of China	China	SCImago #77 · THE 51 · QS =132	1
Xi'an Jiaotong University	China	SCImago #58 · THE 201–250 · QS 305	1

Geographic distribution of citing authors

Country	Citing papers
United States	13
China	11
United Kingdom	3
Singapore	2
Canada	1
Italy	1
Israel	1
France	1
Hong Kong	1
India	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.

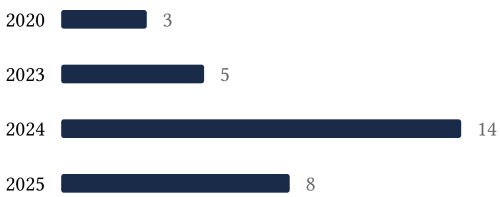

F. AAO Precedent Considerations

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	The Unreasonable Effectiveness of Deep Features as a Perceptual Metric	7	8 CFR 204.5(i)(3) — Outstanding Researcher
Contribution 2	Toward Multimodal Image-to-Image Translation	4	8 CFR 204.5(i)(3) — Outstanding Researcher
Contribution 3	CNN-generated images are surprisingly easy to spot... for now	2	8 CFR 204.5(i)(3) — Outstanding Researcher