

Citation Evidence Report

EB-1B Petition — Outstanding Professor or Researcher

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

Aditya Khosla

CTO, Iterative Health

[Google Scholar profile](#)

Generated 2026-06-10 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

801 Citing papers mapped	801 Citation edges	30 Home papers mapped	38 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.

97.1% independent of 792 classified citing papers

Citation type	Count
Independent	769
Self-citation	3
Co-author	20
Same-institution	0

9 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 foundational benchmarks for large-scale visual recognition and advanced discriminative localization and scene understanding through highly cited, independent work.

The researcher's contribution centers on advancing large-scale visual recognition, anchored by a seminal 2015 paper in the International Journal of Computer Vision. This core work appears to have set a standard for evaluating visual recognition systems on a massive scale, serving as a critical reference point for the field.

Building on this foundation, the researcher's subsequent publications suggest a strategic expansion into more nuanced visual tasks. The 2016 work on discriminative localization indicates an effort to refine how models identify specific regions within images, while the 2017 publication on a ten-million-image database for scene recognition points to a broader focus on contextual understanding and large-scale data infrastructure. This progression implies a move from general classification benchmarks to specialized localization and scene-level analysis.

The significance of this line of work is evidenced by its extensive uptake within the scientific community. The core paper has accumulated over 55,000 citations, while the follow-up studies have garnered over 15,000 and 6,000 citations respectively. Crucially, analysis of citing literature reveals that 97.1% of citations originate from independent researchers, demonstrating that this work has become a widely adopted standard beyond the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 769 · 144 flagged influential by Semantic Scholar

CORE PAPER

[Imagenet large scale visual recognition challenge](#)

2015 · International journal of computer vision 115 (3), 211-252, 2015 · 55,684 citations (GS)

Field-normalised: 42,393 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2015 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Simam: A simple, parameter-free attention module for convolutional neural networks	Shanghai Jiao Tong University, Sun Yat-sen University, The Hong Kong Polytechnic University	China	Methodology
2	Lvlm-ehub: A comprehensive evaluation benchmark for large vision-language models	Georgetown University, Peking University, Shanghai AI Laboratory	China, Hong Kong	—
3	Scaling rectified flow transformers for high-resolution image synthesis	Heidelberg University, LMU Munich, Stability AI	Germany, United Kingdom, United States	Methodology
4	TransUNet: Rethinking the U-Net architecture design for medical image segmentation through the lens of transformers	Alibaba Group (United States), ByteDance, East China Normal University	Australia, China, United States	—
5	Survey on deep learning with class imbalance	Florida Atlantic University	United States	—
6	Synthetic data from diffusion models improves imagenet classification	Google, Google DeepMind, Google Research	United States	Methodology
7	Deep high-resolution representation learning for human pose estimation	Baidu, Microsoft	China, United States	Methodology

No.	Citing paper	Citing institution(s)	Country	S2
8	Depth anything v2	ByteDance, The University of Hong Kong, TikTok	Australia, China, Hong Kong	—
9	Cambrian-1: A fully open, vision-centric exploration of multimodal llms	Korea Advanced Institute of Science and Technology (KAIST), Nanyang Technological University, New York University	Singapore, South Korea, United States	Methodology
10	Depth anything: Unleashing the power of large-scale unlabeled data	ByteDance, The Chinese University of Hong Kong, Zhejiang University, The University of Hong Kong	China, Hong Kong	—
11	A comprehensive review of yolo architectures in computer vision: From yolov1 to yolov8 and yolo-nas	Autonomous University of Queretaro, Instituto Politecnico Nacional, Universidad Autónoma de Querétaro	Mexico	Methodology
12	Object detection in 20 years: A survey	Beihang University, Carleton University, Nanyang Technological University	Canada, China, Singapore	—
13	A comprehensive review of convolutional neural networks for defect detection in industrial applications	University of Huddersfield	United Kingdom	—
14	Res2net: A new multi-scale backbone architecture	Harvard Medical School, Nankai University, Oxford University	China, United Kingdom, United States	—
15	Deep learning for generic object detection: A survey	Arizona State University, National University of Defense Technology, Peking University	Canada, China, Finland	—
16	A benchmark dataset and evaluation methodology for video object segmentation	Disney Research, ETH Zurich, Google Research	Switzerland, United States	Background
17	Deep learning modelling techniques: current progress, applications, advantages, and challenges	Asian University for Women	Bangladesh	—
18	Analyzing and improving the image quality of stylegan	Aalto University, NVIDIA	Finland, United States	—
19	Federated learning for generalization, robustness, fairness: A survey and benchmark	Hong Kong University of Science and Technology, Wuhan University	China, Hong Kong	—
20	A comprehensive survey of few-shot learning: Evolution, applications, challenges, and opportunities	East China Normal University, Institute of Technical Education and Research, Siksha 'O' Anusandhan University, Macau University of Science and Technology	China, India, United States	—
21	Class-incremental learning: A survey	Nanjing University, Nanyang Technological University	China, Singapore	Methodology

No.	Citing paper	Citing institution(s)	Country	S2
22	Cutmix: Regularization strategy to train strong classifiers with localizable features	Naver, NAVER Corp., University of Tübingen	France, Germany, South Korea	Methodology
23	Class-balanced loss based on effective number of samples	Alphabet Inc., Cornell University, Google Research	United States	Methodology
24	Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks	Thomas Jefferson University Hospital	—	—
25	Lvis: A dataset for large vocabulary instance segmentation	Facebook AI Research, Stanford University	United States	Background
26	The cityscapes dataset for semantic urban scene understanding	Daimler AG, Max Planck Institute for Informatics, Mercedes-Benz AG	Germany, United States	Background
27	Image inpainting for irregular holes using partial convolutions	NVIDIA, NVIDIA Corporation	United States	Methodology
28	Generative image inpainting with contextual attention	Adobe Research, ByteDance, University of Illinois at Urbana-Champaign	United States	Result
29	Image-to-image translation with conditional adversarial networks	Adobe Inc., Humen, Inc., Massachusetts Institute of Technology	United States	Methodology
30	Ntire 2017 challenge on single image super-resolution: Dataset and study	ETH Zürich, University of Würzburg	Germany, Switzerland	Background

Showing the 30 most-cited of 769 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).

Citing-text excerpts — how the field used this work

METHODOLOGY Simam: A simple, parameter-free attention module for convolutional neural networks

“All networks are pre-trained on ImageNet-1K (Russakovsky et al., 2015) and transferred to the CO-CO (Lin et al., 2014) dataset by fine-tuning.”

METHODOLOGY Scaling rectified flow transformers for high-resolution image synthesis

“As a widely used dataset, we convert the ImageNet dataset (Russakovsky et al., 2014) into a dataset suitable for text-to-image models by adding captions of the form “a photo of a < class name >” to images, where < class name > is randomly chosen from one of the provided names for the image's class...”

METHODOLOGY Synthetic data from diffusion models improves imagenet classification

“Since the images of ImageNet-1K dataset vary in dimensions and resolution with the average image resolution of 469×387 [48], we examine synthetic data generation at different resolutions, including 64×64 , 256×256 , and 1024×1024 .”

METHODOLOGY Deep high-resolution representation learning for human pose estimation

“In the paper, shown in Table 9. Results on the ImageNet Validation Set We apply our networks to image classification task. The models are trained and evaluated on the ImageNet 2013 classification dataset [54]. We train our models for 100 epochs with a batch size of 256. The initial learning rate is set to 0.1 and is reduced by 10 times at epoch 30, 60 and 90. Our models can achieve comparable performance”

METHODOLOGY Cambrian-1: A fully open, vision-centric exploration of multimodal llms

“ImageNet-1K[105] LanguageSupervised CLIP[102] SSL-Contrastive DINOv2[96] SSL-Masking MAE[45] Diffusion StableDiffusion[104] DepthSupervised MiDaS[13] SegmentationSupervised SAM[61] Figure 2 | Examples of various vision models, objectives, and architectures studied.”

FOLLOW-UP WORK

Learning deep features for discriminative localization

2016 · Proceedings of the IEEE conference on computer vision and pattern ..., 2016 · 15,605 citations (GS)

Field-normalised: 10,483 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2016 indexed by Semantic Scholar, by citation count.

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

FOLLOW-UP WORK

Places: A 10 million image database for scene recognition

2017 · IEEE transactions on pattern analysis and machine intelligence 40 (6), 1452-1464, 2017 · 6,168 citations (GS)

Field-normalised: 4,730 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.

Contribution 2

Claim – Contribution 2

The researcher developed an integrated machine learning framework for stroke prediction, establishing a foundational approach that has been widely adopted by the independent scientific community.

The researcher's contribution centers on the development of an integrated machine learning approach to stroke prediction, as detailed in their seminal 2010 paper published in the Proceedings of the 16th ACM SIGKDD International Conference. This work represents a concrete methodological advancement in applying computational techniques to critical medical diagnostics.

This line of work appears to address the challenge of improving stroke prediction accuracy through integrated machine learning methods. By publishing in a top-tier data mining venue, the researcher introduced a novel framework that likely combined multiple data sources or algorithmic strategies, distinguishing it from prior single-method approaches. The absence of follow-up papers by the same author suggests this core publication stands as a definitive, self-contained contribution to the field.

The significance of this work is evidenced by its substantial citation count of 288, indicating strong uptake within the research community. Furthermore, analysis of citing literature reveals that 97.1% of citations originate from independent researchers, demonstrating that the contribution has had a broad, field-wide impact beyond the researcher's immediate institutional or collaborative network.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 0

CORE PAPER

An integrated machine learning approach to stroke prediction

2010 · Proceedings of the 16th ACM SIGKDD international conference on Knowledge ..., 2010 · 288 citations (GS)

Field-normalised: 211 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2010 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 a foundational framework for multimodal deep learning, a seminal contribution that has been widely adopted and cited by the global machine learning community.

The researcher's core contribution rests on the 2011 paper 'Multimodal deep learning,' published in the Proceedings of the 28th International Conference on Machine Learning (ICML). This work appears to represent a pivotal early effort to integrate multiple data modalities within deep learning architectures, a topic that has since become central to the field. Given the absence of follow-up papers by the same researcher in this specific dataset, the core paper stands as a singular, high-impact artifact of their early influence on multimodal integration techniques.

The originality of this line of work likely lies in its timely proposal of methods to handle heterogeneous data sources, addressing a significant gap in early deep learning research that predominantly focused on single-modal inputs. The title suggests a broad, foundational approach rather than a narrow application, indicating that the researcher aimed to provide a generalizable solution for combining visual, textual, or other data types. This conceptual leap appears to have resonated strongly with the community, as evidenced by the paper's substantial citation count of 5028, which underscores its role as a key reference point for subsequent advancements in the domain.

The significance of this contribution is further highlighted by the high degree of independent uptake. Among 792 classified citing papers, 769 (97.1%) originate from independent researchers, excluding the scholar, co-authors, or same-institution colleagues. This overwhelming majority of independent citations suggests that the work has been widely recognized and utilized by the broader scientific community as a standard or foundational reference, rather than being driven by internal or collaborative citation practices. Such broad, independent adoption is a strong indicator of the work's lasting impact and relevance to the field of machine learning.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 0

CORE PAPER

[Multimodal deep learning](#)

2011 · Proceedings of the 28th International Conference on Machine Learning (ICML), 2011 · 5,028 citations (GS)

Field-normalised: 3,475 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2011 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
Google Research	United States	—	65
Stanford University	United States	SCImago #18 · THE =5 · QS 3	44
Google	United States	—	42
Facebook AI Research	United States	—	37
Google DeepMind	United States	SCImago #90	37
Carnegie Mellon University	United States	SCImago #266 · THE 24 · QS 52	30
Massachusetts Institute of Technology	United States	SCImago #41 · THE 2 · QS 1	30
UC Berkeley	United States	—	27
University of California, Irvine Medical Center	United States	—	27
The Chinese University of Hong Kong	Hong Kong	SCImago #163 · THE =41 · QS =32	26
Meta	United States	—	25
New York University	United States	SCImago #116 · THE =31 · QS 55	24
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	23
DeepMind	United Kingdom	SCImago #90	20

Institution	Country	World ranking	Citing papers
Meta AI	United States	—	20

Geographic distribution of citing authors

Country	Citing papers
United States	447
China	248
United Kingdom	105
Germany	50
Switzerland	40
Canada	39
Singapore	38
Australia	37
France	36
South Korea	29
Hong Kong	23
India	18

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	Imagenet large scale visual recognition challenge	769	8 CFR 204.5(i)(3) – Outstanding Researcher
Contribution 2	An integrated machine learning approach to stroke prediction	0	8 CFR 204.5(i)(3) – Outstanding Researcher
Contribution 3	Multimodal deep learning	0	8 CFR 204.5(i)(3) – Outstanding Researcher